

Learning Network Services for Professional Development

Rob Koper
Editor

Learning Network Services for Professional Development

 Springer

Editor

Rob Koper
Open Universiteit Heerlen
Valkenburgerweg 177
6419 AT Heerlen
Netherlands
rob.koper@ou.nl

ISBN 978-3-642-00977-8 e-ISBN 978-3-642-00978-5

DOI 10.1007/978-3-642-00978-5

Springer Dordrecht Heidelberg London New York

Library of Congress Control Number: 2009926848

ACM Computing Classification (1998): K.3, J.4, H.1, H.3

© Springer-Verlag Berlin Heidelberg 2009

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Cover design: KünkelLopka, GmbH

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Contents

Introduction

1 Introduction	3
Rob Koper	

Section I Social Interaction in Learning Networks

Section Editor: Peter Sloep

2 From Lurker to Active Participant	17
Peter Sloep and Liesbeth Kester	
3 Guidelines to Foster Interaction in Online Communities	27
Adriana Berlanga, Ellen Rusman, Marlies Bitter-Rijkema, and Peter Sloep	
4 Knowledge Dating and Knowledge Sharing in Ad-Hoc Transient Communities	43
Liesbeth Kester and Peter Sloep	
5 How to Trigger Emergence and Self-Organisation in Learning Networks	57
Francis Brouns, Sibren Fetter, and Peter van Rosmalen	

Section II Navigation Services for Learning Networks

Section Editor: Hans Hummel

6 Individualised Navigation Services in Learning Networks	75
Hans Hummel, Bert van den Berg, Adriana Berlanga, Hendrik Drachsler, José Janssen, Rob Nadolski, and Rob Koper	
7 Evaluating the Effectiveness of Personalised Recommender Systems in Learning Networks	95
Hendrik Drachsler, Hans Hummel, Bert van den Berg, Jannes Eshuis, Wim Waterink, Rob Nadolski, Adriana Berlanga, Nanda Boers, and Rob Koper	

8	How to Set Up Simulations for Designing Light-Weight Personalised Recommender Systems	115
	Rob Nadolski, Bert van den Berg, Adriana Berlanga, Hans Hummel, Hendrik Drachslers, Rob Koper, and Peter Sloep	
9	How to Find and Follow Suitable Learning Paths	151
	José Janssen, Adriana Berlanga, and Rob Koper	
Section III Assessment and Placement Services in Learning Networks		
Section Editor: Jan van Bruggen		
10	A Process-Oriented Approach to Support Multi-Role Multi-Stage E-Assessment: A Case Study	171
	Yongwu Miao and Henry Hermans	
11	Placement Services for Learning Networks	195
	Jan van Bruggen, Marco Kalz, and Desirée Joosten-Ten Brinke	
12	Tools and Techniques for Placement Experiments	209
	Wim van der Vegt, Marco Kalz, Bas Giesbers, Fridolin Wild, and Jan van Bruggen	
13	A Validation Scenario for a Placement Service in Learning Networks	225
	Marco Kalz, Jan van Bruggen, Bas Giesbers, Ellen Rusman, Jannes Eshuis, and Wim Waterink	
Section IV Contextualized Learning Network Services		
Section Editor: Marcus Specht		
14	Towards Contextualized Learning Services	241
	Marcus Specht	
15	Mobile Social Software to Support Authenticity	255
	Tim de Jong and Marcus Specht	
16	Supporting Authentic Learning Contexts Beyond Classroom Walls	273
	Jan Herrington, Marcus Specht, Gwyn Brickell, and Barry Harper	
17	Awareness and Reflection in Mobile Learning Support	289
	Christian Glahn, Dirk Börner, Jeroen Storm, and Marcus Specht	
Section V Learning Networks Integrated		
Section Editor: Hubert Vogten		
18	A Conceptual Model of Learning Networks	313
	Rob Koper	

19 The TENCompetence Infrastructure: A Learning Network Implementation 329
Hubert Vogten, Harrie Martens, and Ruud Lemmers

Section VI Implementation Examples
Section Editor: Wolfgang Greller

20 The Personal Competence Manager 361
Hubert Vogten and Harrie Martens

21 Personal Development Planner 369
Hubert Vogten and Harrie Martens

22 Learning Design Authoring Tools in the TENCompetence Project 379
Phillip Beauvoir, David Griffiths, and Paul Sharples

23 IMS QTI Authoring 389
Yongwu Miao

Conclusion of the Book

24 Conclusion of the Book 399
Rob Koper, Peter Sloep, Hans Hummel, Hubert Vogten,
Jan van Bruggen, Marcus Specht, and Wolfgang Greller

Index 411

Contributors

Phillip Beauvoir Institute for Educational Cybernetics, University of Bolton, Deane Road, Bolton BL3 5AB, UK, p.beauvoir@dadabeatnik.com

Adriana Berlanga Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, adriana.berlanga@ou.nl

Marlies Bitter-Rijkema Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, marlies.bitter@ou.nl

Nanda Boers Psychology Department, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, nanda.boers@ou.nl

Dirk Börner Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, dirk.boerner@ou.nl

Gwyn Brickell Faculty of Education, ICT Research Institute, University of Wollongong, Northfields Avenue, Wollongong, NSW 2522, Australia, gbrickel@uow.edu.au

Desirée Joosten-Ten Brinke Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, desiree.joosten-tenbrinke@ou.nl

Francis Brouns Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, francis.brouns@ou.nl

Tim de Jong Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, tim.dejong@ou.nl

Hendrik Drachslér Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, hendrik.drachslér@ou.nl

Jannes Eshuis Psychology Department, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, jannes.eshuis@ou.nl

Sibren Fetter Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, sibren.fetter@ou.nl

Bas Giesbers Faculty of Economics and Business Administration, University of Maastricht, Tongersestraat 53, 6200 MD Maastricht, The Netherlands, s.giesbers@maastrichtuniversity.nl

Christian Glahn Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, christian.glahn@ou.nl

Wolfgang Greller Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, wolfgang.greller@ou.nl

David Griffiths Institute for Educational Cybernetics, University of Bolton, Deane Road, Bolton BL3 5AB, UK, dai.griffiths.1@gmail.com

Barry Harper Faculty of Education, ICT Research Institute, University of Wollongong, Northfields Avenue, Wollongong, NSW 2522, Australia, bharper@uow.edu.au

Henry Hermans Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, henry.hermans@ou.nl

Jan Herrington Faculty of Education, ICT Research Institute, University of Wollongong, Northfields Avenue, Wollongong, NSW 2522, Australia, janherrington@gmail.com

Hans Hummel Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, hans.hummel@ou.nl

José Janssen Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, jose.janssen@ou.nl

Marco Kalz Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, marco.kalz@ou.nl

Liesbeth Kester Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, liesbeth.kester@ou.nl

Rob Koper Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, rob.koper@ou.nl

Ruud Lemmers Logica Nederland B.V., Robert Schumann domein 4, 6229 ES Maastricht, The Netherlands, ruud.lemmers@logica.com

Harrie Martens Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, harrie.martens@ou.nl

Yongwu Miao Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, yongwu.miao@ou.nl

Rob Nadolski Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, rob.nadolski@ou.nl

Ellen Rusman Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, ellen.rusman@ou.nl

Paul Sharples Institute for Educational Cybernetics, University of Bolton, Deane Road, Bolton BL3 5AB, UK, p.sharples@bolton.ac.uk

Peter Sloep Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, peter.sloep@ou.nl

Marcus Specht Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, marcus.specht@ou.nl

Jeroen Storm Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, jeroen.storm@ou.nl

Jan van Bruggen Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, jan.vanbruggen@ou.nl

Bert van den Berg Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, bert.vandenberg@ou.nl

Wim van der Vegt Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, wim.vandervegt@ou.nl

Peter van Rosmalen Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, peter.vanrosmalen@ou.nl

Hubert Vogten Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, hubert.vogten@ou.nl

Wim Waterink Psychology Department, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands, wim.waterink@ou.nl

Fridolin Wild Vienna University of Economics and Business Administration, Vienna, Austria, f.wild@open.ac.uk

Introduction

Chapter 1

Introduction

Rob Koper

1.1 Rationale of the Book

In 2003 we started a new research programme at the Centre for Learning Sciences and Technologies (CELSTEC) that was aiming to help people to further develop their professional competences by using the innovative powers of new media, mobile devices, and modern Internet services. The idea behind the programme was to contribute to one of the bigger challenges in our society: how to deal with the growing complexity, the growing quantity and the permanent changes in knowledge and technologies. For companies this question relates to the core of their business: how to become innovative and stay competitive. For the employees, the ‘professionals’, this question relates directly to their jobs: how to become and stay employable. In this book we will concentrate on the last group, the professionals and their question how to stay employable, how to keep up-to-date and how to develop professional competence during their careers. The professionals represent the human capital, the knowledge, the innovative power in our economy. The challenge of permanently changing knowledge and technologies puts an enormous pressure on individuals to keep-up with all these changes in order to be able to function effectively in their work. Most of the current initiatives, especially under the umbrella of ‘lifelong learning programmes’, try to cope with this phenomenon by intensifying the offering of courses and training events that are delivered at a distance, by e-learning or in the classroom. Also schools and universities are opening up their classes to professionals. This is combined with measures to stimulate or even force professionals to follow all these courses and training events, for instance by introducing a system in which you have to earn a certain amount of Continuing Professional Development (CPD) points each year, a system that is for instance

R. Koper (✉)

Centre for Learning Sciences and Technologies, Open University of the Netherlands,
Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands
e-mail: rob.koper@ou.nl

well-known in the medical sector. This is a good first step to solve the problem of continuing professional development, but most of the training offerings are taking the professional outside of the context of their daily work and life for some time, to sit in class rooms and conference venues to be trained. This traditional educational method is not always the most suitable for all types of learning and especially not suitable for every professional, because professionals tend to be very busy people, have their families, their travels, their lives full of other priorities. For many of them, the available time to learn outside of the context of their work and daily life is very limited, or even absent. Also, the transferability of the knowledge gained in these courses to the real practice of the professional is questionable, as everybody knows who has been in these events: most of the time you are working very enthusiastically on a topic for some time during the event, but back in the office it is 'business as usual' and the new approaches are soon forgotten because they do not fit the specific circumstances you work in, or are not applicable in your case. Another issue is that professionals are adults with many different experiences and backgrounds, they differ highly in learning needs, preferences and prior knowledge. So, a method of professional development will only be efficient when it is as adaptive, and personalised as possible. Another factor to mention is that the volume of new knowledge produced worldwide in a profession is so enormous nowadays that the traditional role of the 'teacher as a knowledge source', doesn't work in these highly specialised areas anymore. Before a teacher or trainer is fully educated, the knowledge has already changed. This means that professionals should be given direct access to the source of new knowledge and innovative solutions, and the best way to do this is to involve them directly in the innovation and knowledge production itself as part of their job. Related to this is also the acknowledgment that careers of individuals can vary quite a lot, because jobs and market demands are changing. People tend to have various different jobs, even in different sectors during their professional career. Systems like the CPD points are only valuable within the same professional area and not for the planning and management of individual careers with the possibility of multiple jobs in multiple sectors.

A last argument for an additional approach to professional learning is that the demand for high quality, specialised learning and its enormous diversification in all sorts of professional areas is growing so fast, that we simply are not able to provide and pay enough teaching and training staff (and classroom time for the professional) to keep-up with the actual growth in demand.

On the other side, we are also rather sceptical about the power of self-directed learning as a panacea for all learning needs, especially in domains where the changes are enormous and sources are difficult to find. Just reading news from books, blogs and other news sources in your area is not sufficient to acquire new skills, complex knowledge or to leverage yourself to a higher level of functioning. The availability of help, guidance, planning and direction can be an essential requirement to attain more ambitious learning goals and to make the whole process of learning more efficient. This also includes the possibility to separate yourself sometimes from your

work and go to specific places, like a conference venue, a training facility or even a classroom when this leads to more efficient learning.

So, we have taken it as our challenge to come up with some alternative ways of professional learning, which is taking care of the specific characteristics of the (adult) professional, and mixes the best of both worlds: the more informal self-directed professional development with the more formal type of teacher-directed training. We see it furthermore as our challenge to create a seamless integration of this new type of professional learning into work and daily life and make the learning experience more intensive and more efficient for the professional.

Fortunately, the enormous possibilities of new media, mobile devices and especially the Internet, i.e. Web 2.0 social software services, are very helpful in bringing people in contact with each other and in contact with the direct sources of knowledge. So our idea is to provide network services to connect people, knowledge, training possibilities and ideas in what we call a *Learning Network* and provide various *Learning Network Services* that facilitate the learner to access the Network from everywhere, to navigate through the Network, to position themselves in the Network in terms of competences and knowledge, to find people who can help them with specific questions and to get the support of peers and other professionals. In a Learning Network, professionals are facilitated to:

- exchange experience and knowledge,
- collaborate in joint innovations and research,
- offer and get support for further informal, non-formal and formal professional development in the field,
- monitor and identify changes in the field and translate these to changed competence requirements for professionals in the field,
- organise workshops, discussions and conferences,
- to offer support for the assessment of prior learning, including qualifications, certificates and/or diplomas,
- support each other when encountering learning problems,
- use tools and services to register and monitor progress, to create personal development plans, to manage competence profiles and to author learning resources.

These communities are facilitated by Learning Network Services. Specifically the question which Learning Network Services are needed and how they are designed and implemented was the focus of most of our research. We studied and experimented with the Learning Networks concept for five years in various different projects, like Ph.D. Projects, European projects, internal research and development projects and pilot testing projects. This book summarises our experience and knowledge gained in the field of what we call 'Learning Network Services for Professional Development'. In this section you are introduced to the aim of the book, along with a basic understanding of the core concepts of Learning Networks to provide you with a sufficient overview to read this book. Guidelines for reading are also provided.

1.2 Aim of the Book and Target Readers

The aim of the book is to present contemporary, research-based insights into the field of Learning Networks and especially into the web-services that can be provided to facilitate the processes within these networks. The main emphasis of the book is to explain what services a Learning Network requires and what the reader should do to design and run an effective Learning Network Service, including guidelines how to evaluate the effectiveness of a Learning Network Service.

The book is meant to be a worthwhile source of information for practitioners and professional development providers who want to stimulate learning of professionals through social interaction within a company, university, school, region, etc. Because the book is based on research, it is also a good introduction for other researchers in the field who want to be updated in this topic and for students in the field of Learning Sciences or Technology Enhanced Learning. Furthermore, for managers of educational institutions and training companies it is a source of information to provide Learning Network Services to their students and customers, and last but not least the book provides a source of technical information for ICT developers who have to programme specific Learning Network Services.

Most chapters do not require any specific pre-knowledge to read the book, although the book expects a basic understanding of the issues in professional development and technologies like e-mail, Web services, Web 2.0 and Mobile technologies. Some chapters are intended for designers and programmers who want to setup new Learning Network Services, these chapters require more extensive programming and design experience. They will be identified in the section 'How to Read This Book'.

1.3 Learning Networks

A *Learning Network* is defined in this book as a technology supported community of people who are helping each other to better understand and handle certain events and concepts in work or life. As a result, participating in Learning Networks stimulates professional development, a better understanding of concepts and events, career development and employability. Participating in Learning Networks can be a worthwhile instrument for learning, alongside the more regular, formal forms of education that we all know (they can even be integrated into it). Examples of Learning Networks are:

- The employees of a company who want to learn how to provide customer services for new products.
- Communities of teachers who exchange their experience in handling certain pedagogical situations in the classroom.
- Parents of disabled children who exchange their experience and help each other to learn how to handle certain situations in the best way.

- Communities of researchers who exchange information to find solutions about a pollution problem. They update each other with new knowledge and solve problems jointly.
- Lawyers who exchange knowledge and experience when a new law is introduced within their field.
- Students who help each other to write a dissertation.
- A football, golf, tennis or other sports club (mixes competition = assessment with informal and formal training opportunities and also knowledge exchange).

Most of the current Learning Networks are still very weakly supported with information and communication facilities. The intention is to enrich these existing Learning Networks with Learning Network Services that makes them more efficient and makes them better accessible from everywhere in order to integrate them into daily life and work.

Learning Network Services are Web-services that are designed to facilitate the members of the network to exchange knowledge and experience in an effective way, to stimulate active and secure participation within the network, to develop and assess the competences of the members, to find relevant peers and experts to support you with certain problems, and to facilitate ubiquitous and mobile access to the Learning Network. Figure 1.1 provides a concept diagram that contains most of the key components of a Learning Network in a professional context. A more elaborated and generic model can be found in Chap. 18.

In this diagram (me) is positioned with all the relationships within a work context: colleagues at work and other people in the same profession. Not drawn, but existing are all other relationships at home, sports, online, etc. A Learning Network is build around the networks of individual persons, just like this is done in social software environments like elgg, facebook, hyves, linkedIn, etc. So, Learning Networks are organic, they are self-organised by nature.

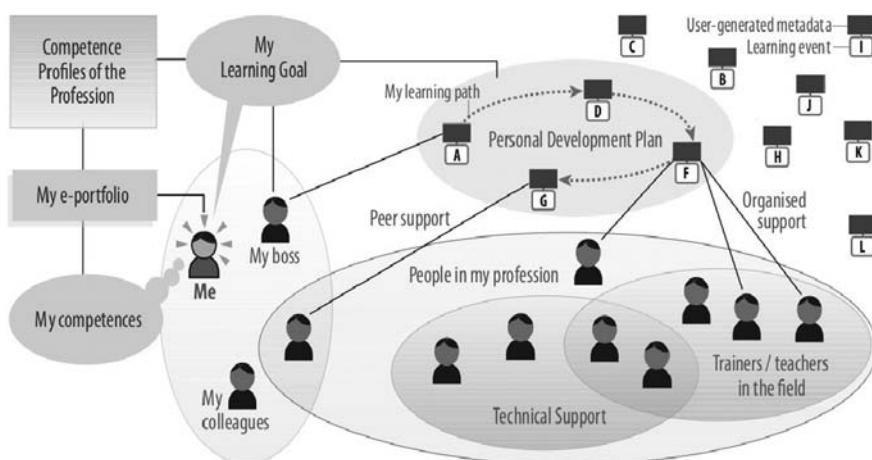


Fig. 1.1 Key components of a Learning Network of a professional

Within a Learning Network a person can define learning goals that could be triggered by the person self or externally (e.g., by *my boss*). Examples of such learning goals are, that you want to be updated on the latest developments in a specific area, want to upgrade your competences to work towards a better or different job or function, want to improve the proficiency level of a specific competence. A learning goal is formulated within the context of formal or informal defined competence profiles that specify the minimal required levels to function adequately, e.g., a job or function profile, but also more informally some statements or opinions about what is needed to be a good professional. The levels of the competences the person actually has are (self-)assessed and compared with the required levels as specified in the competence profiles. The result is a kind of gap analysis per competence that is stored in my e-portfolio and that provides the basis for the creation of a personal development plan (PDP). A PDP consists of a series of learning events that represent education and training activities like courses and training events, but it can also be an expressed intention to look after some behavioural or knowledge aspects (Improve my presentation skills). This intention is only translated to a learning event when the occasion occurs that the person has to do a presentation in the work context. The results are then reported later, e.g. in a blog.

The learning events can be sequenced and planned towards an optimal learning path that the person intends to follow in the PDP. Each learning event has some user-generated metadata connected to it as soon as the event has been performed by other professionals earlier. Examples of this kind of metadata are for instance: what is the rated quality of the event, with what rate of success did persons follow the event, what prior knowledge did the successful participants of the event had, etc. This information provides a valuable base for recommender systems that recommend learning events based on the experience of others.

When someone encounters a learning event, e.g., is busy solving some problem in the work context, one can have a need for the support of another person. So, facilities will be available to find peers or trainers/trainers/coaches who can help you at a specific moment.

Given this initial conceptional model of a Learning Network, we can now identify the various Learning Network Services that are described in this book.

In Section I the focus is on services that stimulate social interaction, e.g., finding suitable peers to help you solving problems, and furthermore incentive mechanisms to stimulate persons to help others.

In Section II the focus is on navigation services that recommend adequate learning events, based on the user-generated metadata.

In Section III the focus is on the gap analysis, i.e., the assessment of competence levels, given certain competence profiles and the selection of the related set of learning events that can be picked and mixed from in the PDP.

In Section IV the focus is on the mobile access of the Learning Network Services and the personalisation and contextualisation of the learning events itself.

In Section V all these aspects are integrated into a common model for Learning Networks, providing the basis for technical Learning Network Services and tools.

In Section VI the focus is on concrete implementation examples.

1.4 How to Read the Book

This book deals with various Learning Network Services that each are described in a separate section of the book. Table 1.1 represents the structure of the book. It presents the topic of the sections, the section editors and the primary reader groups for the section (in priority order).

Table 1.1 Structure of the book

Learning networks service	Section	Section editor	Reader groups
Services that support social interaction in Learning Networks	I	Peter Sloep	1. Providers 2. Researchers
Services that support the navigation within Learning Networks	II	Hans Hummel	1. Providers 2. Researchers
Services that support the assessment and placement of persons within a competence framework	III	Jan van Bruggen	1. Providers 2. Researchers 3. ICT developers
Services that provide adaptation in and mobile access to a Learning Network	IV	Marcus Specht	1. Providers 2. Researchers 3. ICT developers
Infrastructure services that are needed to create an integrated Learning Network	V	Hubert Vogten	1. ICT Developers 2. Researchers 3. Providers (first chapter only)
Implementation examples	VI	Wolfgang Greller	1. ICT Developers 2. Researchers

Each section has a Section ‘Introduction’ that provides an introduction into the section. It explains the type of Learning Network Service that is addressed in the section and explains how it helps users within Learning Networks to do their tasks more effective or efficient. Furthermore it summarises each chapter in the section and provides information about the way different reader groups (providers, researchers and ICT developers) can best read the chapters in the section.

The last chapter of this book is a ‘Conclusion’ where we reflect back on the original goals of the research as it is described in this introduction. What knowledge did we gain exactly? Which knowledge can be applied in practice and which should be research in more detail? Which areas are still underdeveloped and should have a practical or research focus in the coming years?

At the beginning of each section and chapter we have generated a *tag cloud* that provides a worthwhile first impression of the concepts that are addressed in the chapter or section. The tag cloud of the whole book can be found in Fig. 1.2.

- Sydney Olympic Park Authority funded the development of the Sydney Olympic Park project in collaboration with NSW Department of Education and Training Curriculum Directorate and Paramatta Diocese Curriculum Unit.
- The Environmental Science Experiences for Schools (Killalea) project was supported by the Department of Education, Science and Training through their Australian School Innovation in Science, Technology and Mathematics program, ASISTM. The project partners consisted of twelve local schools, the NSW Department of Education, the University of Wollongong, CSIRO Entomology, NSW Department of Environment and Conservation National Parks Division and NSW Department of Lands- Killalea State Park Board of Trustees.

Last but not least, we want to thank Sabine Roks for her support in the final editing process.

not rely on the kind of one-stop solutions that present-day schools and universities provide, nor does it necessarily rely on fixed curricula, classroom instruction, and cohort-based pacing. Non-formal learning takes the desires of 'students' as its starting point rather than institutional offerings (Colley et al. 2003; Schugurensky 2000). It thus is pull-based rather than push-based, if you like (Sloep et al. 2008).

Thus far, we have characterised non-formal learning by exclusion, by describing what it is not and does not assume. This prompts the question of how non-formal learning may become a reality. If institutions such as schools and universities with their lecture rooms and curricula are not the answer, what is? It is our claim that Learning Networks are the devices that should come in their place. This section discusses the social aspects of Learning Networks will be discussed. First, the question will be addressed *why* non-formal learners would bother to act socially (Chap. 2). What is in it for them? Given their busy lives, perhaps having to fit learning into a schedule filled with work, family, and leisure obligations, this is a valid question. And if indeed it is useful for them to engage socially in a Learning Network, how then can they be convinced of this? Second, the question of *how* sociability in Learning Networks best could emerge and be maintained will be taken up (Chaps. 3, 4 and 5). Chapter 3 looks at guidelines for the maintenance of the patchwork of communities that will arise within the boundaries of a specific Learning Network, Chap. 4 discusses in detail guidelines that should guarantee the emergence of such communities. For this a new notion is introduced, that of ad-hoc transient communities. Such communities provide the mechanism for community emergence, the argument is. Chapter 5 describes a case in which they can be seen in action.

At this juncture, a word of caution is in order. The above characterisation of non-formal learning may seem to indicate that thinking in terms of Learning Networks has no bearing on formal learning at all, perhaps even seeks to ban it entirely from the landscape of education. That would be a grave mistake for at least two reasons. First, in all likelihood the initial education of children and adolescents will be best served by a formal approach to it, even if reforms may be in order. Indeed, in formal education, particularly in vocational formal learning, attempts are being made to adapt the traditional push model and make it adopt features of the kind of pull model we advocate here (Anonymous 2007).

Second, there is no reason why, in the context of Learning Network, bouts of formal learning could not be incorporated if those happen to be the most efficient and effective way for particular learners to cater for their competence needs. The reason why formal learning is downplayed in this section is because much of our current expertise in schools and universities is with the push model. So promoting a pull model requires a rethinking of much conventional wisdom. This pertains to many of our traditional educational assumptions, but also to the organisational aspects of the educational universe that is needed, and to the business models that underpin the economic viability of such a universe. Thinking in terms of Learning Networks allows us to break away from conventional wisdom, precisely because several of the traditional assumptions that one surreptitiously makes, are abandoned or at least questioned. Indeed, it the unconventional attitude which thinking in terms of Learning Networks requires that may teach us valuable lessons for formal learning as well.

References

- Anonymous: *Flexibiliteit als voorwaarde; onderwijslogistiek in vraaggestuurd MBO* (Stichting Kennisnet ICT op school, Zoetermeer 2007) p. 24
- Colley, H. et al.: *Informality and Formality in Learning: A Report for the Learning and Skills Research Centre* (Learning and Skills Research Centre, London 2003)
- Communities, C.o.t.E.: *Commission Staff Working Paper: A Memorandum on Lifelong Learning* (European Commission, Brussels, Belgium 2000)
- Edwards, R., Usher, R.: Lifelong learning: A postmodern condition of education? *Adult Educ. Q.* **51**(4), 273–287. (2001)
- Griffin, C.: *Curriculum Theory in Adult and Lifelong Education* (Groom helm, Kent 1983)
- Longworth, N., Davies, W.K.: *Lifelong Learning, New Vision, New Implications, New Roles for People, Organisations, Nations and Communities in the 21st Century* (Kogan, London 1996)
- Schugurensky, D.: The forms of informal learning: Towards a conceptualization of the field. NALL Working Paper No. 19 **19** (2000)
- Sloep, P. et al.: A European research agenda for lifelong learning. In *EADTU Annual Conference 2008, Poitiers France 18–19 September 2008 Lifelong learning in higher education: networked teaching and learning in a knowledge society.*, Poitiers, France 2008
- Sloep, P., Jochems, W.: De e-lerende burger. In: *Jaarboek ICT en samenleving 2007; gewoon digitaal*, ed. by Steyaert, J., De Haan, J. (Boom, Amsterdam 2007) pp. 171–187

as a group. However, for the case of the emergence of sociability, we'll take the worst-case scenario of a collection of unconnected individuals. If sociability can be made to emerge in an environment that resembles a social desert, it always will, is the argument.

This implies that a Learning Network may not be equated with a kind of community. Over time, a Learning Network could develop community-like characteristics. Indeed, we argue that it should in order to maximise its utility. After all, there is much to be gained for the inhabitants of a Learning Network from the mere fact that they all share a particular interest, which is specific to this Network. These benefits may materialise in two different ways, already distinguished in 1973 by Mark Granovetter in a seminal paper (Granovetter 1973). If one thinks of social communities, the mental image is mostly that of a close-knit community in which everybody pretty much knows everybody else. Ties between people in such communities are strong, in the sense that they interact frequently and intensely. Although this is a virtue in that social interactions run smoothly, it is a vice in that knowledge not available inside the community will for ever elude its members. In the words of Granovetter 'strong ties, breeding local cohesion, lead to overall fragmentation' (p. 1378). To access others outside the local community, weak ties have to be exploited, or 'weak ties . . . are . . . indispensable to individuals' opportunities and to their integration into communities' (p. 1378). In line with this observation, it is our central thesis that a Learning Network, being devoid of communities in its incipient phase, provides ample opportunities for community emergence and growth, and hence the establishment of strong ties, through the exploitation of the many weak ties it harbours. In the end, therefore, we view a Learning Network as consisting of many, partly overlapping communities. Through the communities, the benefits of a strong-knit community are reaped, through the overlap, information may flow through the Network as a whole (see also Burt 2000; Reagans and McEvily 2003).

In this chapter we will specifically go into the question of how prospective Learning Network users may be convinced of these benefits, for that is likely to be the necessary condition for their active participation in any Learning Network. Their question would be 'Why should I participate?', this chapter inventories answers to that question, which are then translated into a few guidelines for those contemplating to set up a particular, topic-bound Learning Network. Two kinds of answer are distinguished. Proximate answers, which affect the decision to participate here and now; and ultimate answers, which motivate participation, but only in the long run, after the decision to participate has already been taken. Both are important, the former to persuade people to participate, the latter to persuade people to keep participating. Before going into them, we'll introduce a concrete example to add some realism to the discussion.

2.2 The Moto Guzzi V7 Enthusiasts

Eddy LeDuca is 38 years old and recently bought an old Moto Guzzi V7 from 1972. To restore it in its original state and make it operational again, he wants to learn

how to go about that. From a colleague he got wind of an online vintage motorcycle network. In it he hopes to learn some tips and tricks for renovating his newly-bought vintage Guzzi V7.

Jannie Barends is 62 years old and enjoys an early retirement. She bought a brand new Moto Guzzi V7 back in 1972 and now owns twelve motorcycles, all collectors' items. Since the Guzzi was her first motorcycle she is very attached to it and does everything she can to keep it running. She has a whole library of manuals on how to maintain, rebuild and repair motorcycles, and is used to exchange information with other motorcycle fanatics.

Bas Timmer is 23 years and works as a car mechanic at 'Stop and Go', a franchise specialised in small car-reparations that are done while you wait. He has the ambition of running his own garage in the near future. By way of preparation, he surfs the Internet and visits online discussions and fora on cars and motorcycles. He wants to keep his knowledge up to date and stay on top of what lives among the car and motorcycle amateurs. In the business plan for his own garage he wants to include services that match the needs of the amateurs.

Jessica Zwart is 41 years old and works for the research and development department of Moto Guzzi. She is an experienced person. Ever since the advent of the Internet, she became an active member of all kinds Moto Guzzi discussion fora. Like Bas, she uses them to keep informed about what lives among owners of vintage motorcycles, Moto Guzzis in particular. Her regular posts are intended primarily to gauge customer satisfaction, and test new research and development ideas.

Eddy, Jannie, Bas and Jessica all share a passion for vintage motorcycles. In one way or another, Eddy, Jannie, Bas and Jessica are really all lifelong learners who – from their various perspectives – want to expand their knowledge about vintage motorcycles, in particular the Moto Guzzi from 1972. Wouldn't there be a better way to serve their interests than is done currently by rather haphazardly surfing the Internet and, every so often, engaging in a discussion forum? Joining the Learning Network on vintage motorcycles seems to be a good idea, but what would be convincing arguments to them?

2.3 The Long-Term Perspective

Arguments to convince Eddy, Jannie, Bas and Jessica should refer to the ways in which each one of them personally benefits from sharing knowledge with the other participants in the vintage motorcycle Network. For Eddy, this would relate to his ability properly to renovate the bike, for Jannie the ties she develops with fellow-enthusiasts, for Bas the insights he gains in how to set up his own bike shop in due time, and for Jessica the user feedback she receives. These all refer to motives for participation the beneficial effects of which reveal themselves in the long run. Such motives come in a few kinds.

First, note that the reasons why knowledge is exchanged in some Learning Network may range from purely educating oneself, such as done by Eddy and perhaps

Jennie, to developing oneself professionally, such as done by Bas and Jessica. Thus the use of a Learning Network extends beyond the educational realm into the participants' professional life, present and future. A Learning Network qua knowledge sharing community thus acquires characteristics of a community of professionals. Particularly to someone such as Bart this is very significant. While learning about vintage motorcycles, in his case particularly the Moto Guzzi V7, he comes in contact with many people, such as Jessica, who will be useful to him in his future professional life as a bike shop owner. This applies generally. The communities of learning that arise in a Learning Network may acquire characteristics of communities of professionals (Brown 2001). As argued, typically learners in a Learning Network combine their need to learn with the necessity to work. Indeed, their learning needs often derive from their occupation. So there is every reason to expect that the communities that arise in the Learning Network will acquire this dual nature of a learning community and a professional community (Longworth and Davies 1996).

Second, as has been pointed out by Nardi et al. (2000), it has become less productive only to rely on knowledge sources within the company you happen to work with. Such sources have become less reliable and less accessible with the increased turnover rate of personnel and indeed companies themselves. If your company is a constant state of flux and you yourself are in constant danger of being replaced or even losing your job, it is much more productive and sensible to rely on your own, personal network, a type of network she describes as *intensional*. (This kind of network, parenthetically, shows remarkable resemblance with the ad-hoc transient communities discussed in Chaps. 4 and 5) This makes you as a person less dependent on the company you work for. In addition, the chances of accessing novel information are increased since you step out of the probably close-knit group you are part of (Burt 2000; Reagans and McEvily 2003).

Third and focussing specifically on the act of learning itself, there is ample evidence that collaboration and a social setting significantly improve learning effectiveness and learning efficiency. By collaborating with others, learners make use of their collective intelligence, motivate and enlighten each other and thus improve their learning outcomes (Allen 2005; Cartney and Rouse 2006; Chapman and Ramondt 2005; Keppell and Au 2006). Some will say they have become part of a community of learning (Wilson and Ryder 1998). In educational circles, this is a familiar argument, which goes back to the ideas Vygotski (1978) or even Dewey, back in 1916 (Dewey 1916). Related to this but different from it is the argument that helping others in a learning context, that is acting as peer-tutors, is a powerful learning experience in and of itself (Fantuzzo et al. 1989; Wong et al. 2003) (see for more details Chap. 4).

Therefore, we have uncovered two kinds of reasons for participating in a Learning Network: it benefits you as a professional, prospective or actual, and it improves your learning. These benefits materialise in the future as a consequence of your having been active in the Network. Guidelines for Learning Network designers will have to consider this long-term character. They should point out these benefits to the novice users of a Learning Network, perhaps through accounts of successful participation of past and present users. Guidelines thus take the form of information

about and explanations of benefits. These may convince people to ‘give it a try’; they will not convince them to change from passive onlookers into active participants. Although there is a role for such ‘lurkers’, see Preece et al. (2004), a Network of only lurkers will rapidly lose its attractiveness. So how can lurkers be convinced to contribute actively?

2.4 The Short-Term Perspective

To investigate the question of why some Learning Network participant would decide actively to participate in it, consider the following situation. Eddy, having taken apart the fuel system of his Guzzi V7, finds out that he is hesitant about the exact way in which to reassemble the carburettor. The manual he has shows an exploded view of the carburettor, but his seems to be a slightly different model. Perhaps some of the V7s were fitted with a different model? He decides to seek help. Suppose, a mechanism is in place that allows him to target specific people in the Network who should be knowledgeable about his question. Suppose, Jessica receives his question about how to reassemble his particular make of carburettor. Why Jessica would answer Eddy, is the question. What is in it for her? She might consider to answer Eddy in the hope that next time, when she has a request to Eddy, for instance to gauge his opinion on a new design, he will reciprocate. However, what guarantee does she have he will?

This kind of situation has been analysed extensively in game theory. It is akin to the classical prisoners’ dilemma, in which two prisoners facing a long period of incarceration, have to decide either to stay silent about their misdeed or to confess (Aronson and Thibodeau 1992). If they collaborate and both stay silent, their punishment is smallest (say, each 1 year). If one of them talks and the other does not, the prisoner who talks is worst off (5 years); the prisoner who keeps his mouth shut profits by having his jail time reduced to naught. However, if both talk, they are worst off, as both are sent off for 3 years. The best strategy therefore is to join forces and not talk, however, how can the other person be trusted not to go for no jail time at all by talking? The result of the individually most sensible decision (talk) produces the collectively worst outcome (a total time of 6 years rather than 2). Translated to the example, Jessica should therefore decide not to honour Eddy’s request for help for fear of not being helped by Eddy later on with her request for help. And indeed, what guarantee does she have Eddy will reciprocate?

The predicament can be overcome by repeatedly ‘playing the game’, a situation which is called the iterated prisoners dilemma. The best strategy to follow, simulations have shown, is the tit-for-tat strategy: always cooperate on your first move (help Eddy) and then copy the last move of your opponent (if Eddy failed to reciprocate, Jessica will not help him next time around, if he did, so will she) (Axelrod 1984). The simulations Axelrod carried out for this situation, however, show that a few conditions need to be met for this to work. First, participants need to be identifiable, i.e. have a persistent identity, even if it is a pseudonym. Second, there may be no, to the participants known ending to the ‘game’. If there is, the players do

not have a means of punishing defection behaviour (i.e. cheating on your opponent), so the rules of the one-off prisoners dilemma apply to the last move. However, now they are unsure about their last move, by the same argument they also are about the one but last move, and so on, down to the present move. Third, even though the value of future encounters may decrease relative to the present one – for all you know, there may be no next encounter – the decrease should be limited. Otherwise, if there hardly is a future we are back again at the one-off prisoners dilemma.

It is important to notice that, if the conditions discussed are met, Axelrod's simulations show that cooperation will arise and spread spontaneously in many cases. Even if only a small percentage of a group plays tit-for-tat (about 5%), they can 'invade' a group of people who refuse to collaborate. The upshot is that in an incipient Learning Network in which the conditions just discussed obtain, collaboration will occur. Only some inhabitants will have to be willing to take the risk of answering a question without guarantee of reciprocation. Obviously, guidelines for a Learning Network designer pertain to implementing Axelrod's criteria and keeping the investment needed actually to honour a request for help as low as possible. Chapter 5 discusses an experiment in which this has been done. Although the details will differ from Learning Network instantiation to instantiation, it will not be difficult to prevent people from changing their identities or a Network from ending at a specific date. It is more difficult to have the future cast a sufficiently large shadow into the future, as this has to do with frequency and intensity of contact and size of the Network. No clear-cut guideline may be given therefore.

Although collaboration should arise spontaneously according to a game-theoretical analysis, thus almost pre-empting the need for mechanisms that spur people to collaborate and answer questions, analyses have been made of such mechanisms for different contexts. These hold promises for collaboration in Learning Networks too. We are referring to Stephen Weber's investigation of the mechanisms behind the success of *Open Source* communities (Weber 2004). Like Learning Networks, these communities are in their beginning stages loose-knit and often rely on large numbers of contributors. What motivates them to contribute their source code without any chance of financial recompense, Weber wondered. Of the several mechanisms, he suggests two apply to Learning Network.

First, there is the desire to produce a thing of intrinsic beauty. Although this may be hard to grasp for a non-programmer, it is similar to what one experiences when writing a gripping story, delivering an elegant mathematical proof or cooking an exquisite meal. It has to do with professional pride, something that is hard to experience when producing proprietary software. One should realise that software code is hidden from inspection once the code is compiled, as is necessary to turn it into code that can be executed on an actual computer. Therefore, if the software is proprietary, nobody will ever see it as it is screened off from inspection by trade secrecy, effectively robbing a programmer from the praise from others that may feed his pride and satisfaction as a professional. As open source code can be inspected by

anybody, particularly professional peers, the situation there is completely different. This motive translates to a Learning Network as the pride someone puts into honouring a request for help. Going back to Eddy's carburettor, Jessica could do a quick and dirty answer by pointing to a page in some manual she owns, trusting that Eddy will find a way of accessing it. Or she might write a more elaborate answer, clarifying to him some of the abstruse elements in the manual, a scanned image of which she includes.

Second, now that the code is accessible by everyone, it also provides the programmer an effective means of self-promotion. Anybody, including potential employers and clients, can assess the quality of her work and on the basis thereof decide to hire her. Of course, this motive builds on the previous one. After all, an elegant piece of program code better supports the aim of self-promotion than would a bad instance. This too easily translates to Learning Networks. Jessica has a stake in an answer that is clear and complete. Others who see it will immediately appreciate its quality, something that increases Jessica's reputation as a professional in the Network. And from this, she will profit when she herself needs help with the assessment of the user appreciation of, say, a new design for a Guzzi saddle.

Both examples of incentives that apparently motivate computer programmers to share their code with others will provide incentives to Learning Network participants to honour requests for help. The sense of pride that attaches to having provided an elegant answer and the concomitant benefits to someone's reputation, will obviously only come about if the answer is publicly available. Again, how this translates into guidelines for a specific Learning Network depends on the Network in question. A publicly available, instantly updated list of question asked would be a means. Chapter 5 discusses an experiment with question answering in a Learning Network. In this case, no list of questions asked was provided. However, several people, thus providing a modicum of exposure, discussed questions in wikis. The set-up could however easily have included a list of questions asked and answers given.

A third, rather obvious mechanism to keep in mind, relates to the costs someone incurs who wants to honour a request for help. Such costs consist of two parts. The transactions costs are the effort needed to access the question and process the answer. They need to be added to the material costs of providing the answer itself. Obviously, much is to be gained by keeping the transaction costs low, at the very least their perception. Having to drop a request for help in several, generic fora and then having to check them regularly for an answer, obviously generates high transactions costs. Using RSS feeds to keep a tab on these fora already would lower the costs, etc. Generally, technical solutions will significantly help lower such costs. Perhaps surprisingly, the costs of actually providing an answer can also be lowered by technical means. In the experiment discussed in Chap. 5, participants who have indicated to be willing to answer a content-related question are guided to a wiki, which is seeded with text fragments that pertain to the subject in question. Jointly editing these fragments is a significant reduction of effort compared with thinking up an answer all, individually (Van Rosmalen et al. 2008).

2.5 Conclusion

Learning Network participants will somehow need to be convinced that it is in their interest not to stay lurking and become active participants. Although there are sound ultimate reasons for them to do so – they maximise their profit from the Network both in terms of their professional development and in terms of their learning achievements – they still need to be convinced to honour this specific request from this specific person at this specific moment of time. If the conditions for an iterated prisoners' dilemma apply, collaboration should actually arise spontaneously. However, motives that are more powerful apply. An analysis of the motives that drive programmers to write open source code, revealed two direct motives, the desire to create a thing of beauty and the possibility to contribute to one's professional reputation.

Of course, it remains to be seen whether this suffices. Perhaps elaborate reward and punishment systems are needed. Whatever the case, it is almost certain that what works and does not work crucially depends on the kind of Learning Network to which these guidelines are applied. They will need to receive a local interpretation and instantiation.

In addition, the analysis was deliberately done for first-time users of a novel Learning Network. However, typically, novel users will enter a Network that has been around for a while. Thus, they do not encounter an unstructured whole, but rather a patchwork of communities, which each already address several, slightly different topics within the overall framework that the Network is about. So over time in the vintage motorcycle Network certainly community-like grouping will have emerged that, for instance, as in the case of Moto Guzzi V7. The presence of such a structure will certainly make it easier for first-time users to decide to change their lurking behaviour and become active participants in some specific community. As a Moto Guzzi enthusiast, it is more rewarding to enter into a discussion with fellow enthusiasts Eddy, Jenny, Bas, or Jessica than with someone who is of a different persuasion. How such communities may arise and be maintained is discussed in the two chapters to follow.

References

- Allen, K.: Online learning: Constructivism and conversation as an approach to learning. *Innov. Educ. Teach. Int.* **42**(3), 247–256 (2005)
- Aronson, E., Thibodeau, R.: The Jigsaw classroom: A cooperative strategy for an educational psychology course. In: *Cultural diversity and the schools* ed by Lynch, J. et al. (Palmer, Washington 1992) pp 231–256
- Axelrod, R.: *The evolution of cooperation* (Basic Books, New York 1984)
- Brown, R.E.: The process of community-building in distance learning classes. *J. Async. Learn. Network.* **5**(2), 18–35 (2001)
- Burt, R.S.: The network structure of social capital. In: *Research in organizational behavior* ed by Sutton, R.I., Staw, B.M. (Elsevier, New York 2000) pp 345–423
- Cartney, P., Rouse, A.: The emotional impact of learning in small groups: Highlighting the impact on student progression and retention. *Teach. High. Educ.* **11**(1), 79–91 (2006)

- Chapman, C., Ramondt, L.: Strong community, deep learning: Exploring the link. *Innov. Educ. Train. Int.* **42**(3), 217–230 (2005)
- Dewey, J.: *Democracy and education* (MacMillan, New York 1916)
- Fantuzzo, J.W. et al.: Effects of reciprocal peer tutoring on academic achievement and psychological adjustment: A component analysis. *J. Educ. Psychol.* **81**, 173–177 (1989)
- Granovetter, M.S.: The strenght of weak ties. *Am. J. Sociol.* **78**(6), 1360–1380 (1973)
- Keppell, M., Au, E.: Peer learning and learning-oriented assessment in technology-enhanced environments. *Assess. Eval. High Educ.* **34**(4), 435–464 (2006)
- Longworth, N., Davies, W.K.: *Lifelong learning, new vision, new implications, new roles for people, organisations, nations and communities in the 21st century* (Kogan, London 1996)
- Nardi, B.A. et al.: It's not what you know, it's who you know: Work in the information age. *First Monday* **5**(5), (2000)
- Preece, J. et al.: The top five reasons for lurking: Improving community experience for everyone. *Comput. Human Behav.* **20**, 201–223 (2004)
- Reagans, R., McEvily, B.: Network structure peer and knowledge transfer: The effects of cohesion and range. *Adm. Sci. Q.* **48**, 240–267 (2003)
- Van Rosmalen, P. et al.: A model for online learner support based on selecting appropriate peer tutors. *J. Comput. Assist. Learn.* **24**(6), 483–493 (2008)
- Vygotsky, L.S.: *Mind in society* (Harvard, Cambridge, MA 1978)
- Weber, S.: *The success of open source* (Harvard, Cambridge, MA. 2004)
- Wilson, B., Ryder, M.: Distributed learning communities: An alternative to designed instructional systems. *Educ. Tech. Res. Dev.* (1998)
- Wong, W.K. et al.: Reciprocal tutoring using cognitive tools. *J. Comp. Assist. Learn.* **19**, 416–428 (2003)

Chapter 3

Guidelines to Foster Interaction in Online Communities

Adriana Berlanga, Ellen Rusman, Marlies Bitter-Rijkema, and Peter Sloep



3.1 Introduction

This chapter focuses on the communities that arise within the confines of a Learning Network. As they are not preconfigured, members and all, but are thought to emerge spontaneously, out of their own volition as it were, the question arises how their continued existence may be safe-guarded. What measures and tools, or Learning Network Services, may foster the well-being of communities once they have emerged? How can one make sure that a community, which its members consider valuable, does not disappear because the costs of maintaining it become prohibitive?

As Learning Networks predominantly support non-formal learning in which mainly professional and lifelong learners will participate, the Learning Network Services should take into account their specific characteristics (Koper and Tattersall 2004), in particular in should ensure that learners:

A. Berlanga (✉)

Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands
e-mail: adriana.berlanga@ou.nl

1. are self-directed and take responsibility for their own learning process.
2. can participate, at the same time, in formal and non-formal learning activities.
3. are heterogeneous with respect to competences acquired and sought.

These characteristics imply that online communities for Learning Networks need to be equipped with Learning Network Services to facilitate professional and life-long learners to create and manage themselves. As no institution controls the learning process, they are themselves responsible for their own learning process, and in that respect for organising and directing the network. The services provided by the Learning Network should therefore facilitate and encourage participants to manage the communities they participate in from a bottom-up perspective.

Furthermore, contrary to traditional class-based learning in which all participants more or less have the same level of knowledge, in online communities for Learning Networks the background, competence level and experience of the participants will vary from person to person. They might have different learning goals, working experience, and knowledge about the topic of study but, nevertheless, will have to interact in an online context to meet their individual goals.

To elaborate this further, we will refer back to the case of Eddy LeDuca, introduced in Chap. 2. Apart from being a motorcycle enthusiast, Eddy works as a policy analyst at TsA, *Tomorrow-s-Aqua*, an environmental consultancy firm specialising in water management. Eddy wants to develop his competences in Environmental Sciences, so he can improve his job position with TsA. To improve his competences he has been enrolled for some time in a Master in Environmental Science at the Open Universiteit Nederland, which he is about to finish. He is, therefore, a learner in the particular context of that Master. However, on the topic of water management, Eddy has a lot of knowledge, after all this is his work with TsA. So Eddy's role ranges from being a learner in the context of his Master studies to being an expert in water management in the context of his role as policy-analyst for TsA. Eddy thus participates in at least two different communities, which are part of the same Learning Network on environmental issues. And, needless to say, the opportunities socially to interact with other community members are what make it worthwhile to Eddy to spend time in them. Through social interaction, collaboration, for instance in the form of collaborative learning activities, may come about.

However, in online learning communities social interaction does not happen automatically (Kreijns 2004). For instance, in computer-supported collaborative learning environments the main pitfalls are to take social interaction for granted (because the technology is available learning and social interaction will occur) or the psychological dimension of the act of interacting (because people interact, they will develop trust, a sense of belonging, etc.) (Kreijns et al. 2003). Second, participants who interact online do not have the same kinds of opportunities to learn to know each other as they do in face-to-face situations (e.g., incidental chats, meetings in the corridor, etc.). Consequently, much of the non-verbal signals are missing, impoverishing the communication.

These pitfalls apply generally to online communication. In computer-supported collaborative learning environments they are remedied to some extent by the

presence of an institution which controls the learning process. Learning Networks, in contrast, harbour bottom-up, self-organised online learning communities only that largely rely on the active participation of their members and their willingness to share knowledge. This implies that a third pitfall has to be added, which is mistakenly to assume that communities will flourish merely because participants are registered to the network; that they will be sustainable and prosper in the long term because a central authority (e.g., tutor, teacher, institution, university, etc.) will guarantee so. Such an authority is only available in the very limited sense of maintaining the infrastructure (and perhaps providing a few other services). Fourth, and also contrary to computer-supported collaborative learning environments, Learning Network participants, initially, are not acquainted with each other and do not have a common history. Therefore they do not know who is who in the network, what their specific expertise is and, as a consequence, to whom to turn to for help.

In summary, these four pitfalls make it difficult to build an infrastructure that obeys the necessary conditions for participants to develop their competences. Hence, affordances have to be provided that promote social interaction (Berlanga et al. 2007b) as well as facilitate interpersonal trust formation and an atmosphere of interpersonal trust (Rusman et al. 2007, submitted). In the rest of this chapter, we will explain how to create the social conditions under which people in online learning communities, which are embedded in a larger Learning Network, will sustain their interactions. We will discuss a set of practical guidelines that foster interaction. We start off by describing an example of an informal learning situation: people from different disciplines and backgrounds with a need of developing their competences in the area of environmental sustainability. Just as does Eddy LeDuca.

3.2 The European Environmental Sustainability Community

The *European Environmental Sustainability* community (hereafter EES) seeks to foster a multicultural and multidisciplinary dialogue on sustainable development. This community is organised by a network of European institutions and citizens that share expertise in this area (Cörvers et al. 2007).

Every participant has a different purpose for participation in the EES community. Eddy, for instance, needs to develop his competences in current and new European laws regarding environmental sustainability, so he can get the job position he wants: senior consultant of environmental policies in TsA. Other participants will need to develop their competences on this topic too, but will all have their own reasons. Participants, furthermore, are from different European countries and have diverse backgrounds, such as biologists, chemists, policy-makers, academics; and they work in different companies, institutions, or universities.

This implies that EES participants need to collaborate with each other, find ways to learn together, and help each other. In doing so, they are confronted with different views on sustainability, and different types of expertise from the various disciplinary backgrounds of the participants.

3.3 Guidelines to Foster Sustainability of Online Communities

Sustainability in the context of online communities means tooling online communities in such a way that users may manage, organise, regulate and classify resources, participants and communities. Based on previous work, which analysed best practices in popular and frequently used social web applications (Berlanga et al. 2007b) and explored affordances for knowledge co-construction (Bitter-Rijpkema et al. 2005, 2002), we argue that sustainable online communities should offer services along four dimensions: self-management, self-organisation, self-categorisation, and self-regulation. Each dimension determines a guideline for providing the relevant functionality. Table 3.1 summarises them.

Table 3.1 Guidelines to foster sustainable online communities

Dimension	Guideline
Self-management	Facilitate participants with the creation and management of their own presence as well as their contributions within the community
Self-organisation	Facilitate participants' interaction with others and support knowledge co-construction between them
Self-categorisation	Help participants to classify and evaluate their own contributions but also those from others
Self-regulation	Allow participants to control the level of privacy of (their) contributions, as well as to decide whether these are offensive or not

3.3.1 *Facilitate Participants with the Creation and Management of Their Own Presence as Well as Their Contributions Within the Community*

Social presence theory (Short et al. 1976) argues that the social impact of a communication medium depends on the social presence it allows communicators to have. Computer-mediated environments, such as online learning communities in Learning Networks, do not allow face-to-face communication, non verbal clues are missing and, as a consequence, social presence is weak (Rogers and Lea 2005). To remedy this, participants should be stimulated and enabled to project themselves socially and affectively, to manage their own presence (Rourke et al. 1999). Allowing setting up a profile for themselves and the contacts they gather will help them do so. Allowing them to assemble communities around a topic of their interest and to create learning activities will do so too.

To facilitate participants with the *creation of their profile*, a template should be available. Although developed in a context of formal education, a good example is the pEXPi profile template (abbreviation for *personal expertise inventory* or *personal identity and expertise profile*) (Ogg et al. 2004; Rusman et al. submitted; Rutjens et al. 2003). It is an easy-to-use template that participants of an online community can use to introduce their expertise. The pEXPi template provides a

format that allows members to give context information, communication style preferences and learning ambitions. It serves as a means to start further interaction and enhance the chance of fruitful collaboration. It also provides a personal touch to the (re)presentation of oneself in the online community by offering information such as the participant's picture, name, interests, etc., enhancing a participant's recognisability to others in the community and indeed the Learning Network as a whole. Figure 3.1 shows, by way of example, how our Eddy LeDuca might present himself to the rest of EES participants using the pEXPi template.

It is important to point out that, in our view, *profiles* are not merely a collection of personal data but mainly a means to foster interaction (Berlanga et al. 2008a), encourage participation (Brouns et al. 2007), support initial trust formation between participants (Rusman 2004; Rusman et al. 2007), and promote participants' visibility and awareness of others within the online community (Girgensohn and Lee 2002). The information that the profile should in the first instance contain depends on the goals the community members have. After all, communities in a Learning Network serve no higher order, collective goals; they are merely there to serve the interests of the individuals that make up the community. So, although the inclusion of information such as first name, surname, screen name, and email should be mandatory to guarantee interactions at the network level, participants should be allowed to decide themselves what optional information to display in their profile, as well as to determine the level of privacy attached to it (accessible to the Learning Network as a whole, to specific communities only, to specific people – 'friends' – only, etc.). Optional information could comprise reasons for participating in the online community, preferences, interest, competences to be developed, favourite resources or contacts (Berlanga et al. 2008b).

To support the *creation of contacts* in online learning communities, the services should allow that contacts can include, for instance, peers, teachers, tutors, institutions or even true friends (Boyd 2006). It is desirable, furthermore, to have graphical representations of connection between contacts, this supports awareness, social interaction and the exchange of reputational information (see also Sect. 3.4.2). Having contacts stored using a common format, such as *Friend Of A Friend* (<http://www.foaf-project.org/>), not only makes it easier to create an overview of the members of the Learning Network (Vogten et al. 2008), but also facilitates the exchange and the use of common information about participants' contacts.

Supporting the *creation of online learning communities* in Learning Networks could be done at the level of creating online (sub)communities for formal or non-formal learning. For instance, imagine that Eddy LeDuca would like to create a new non-formal learning community, called *EES courthouse* so EES participants can exchange knowledge and experience about current and new European laws regarding environmental sustainability. It is important that participants can create this kind of (sub) communities, to foster interaction and knowledge sharing.

The last recommendation to be included in this guideline is to support the *creation of learning activities* incorporating, for instance, resources, courses, lessons, assessments, learning materials, and so on. Participants should be able to create their own learning activities, modify existing ones, and distribute them for sharing

pEXPi

Personal expertise page

Personal:		 <p><i>Insert a picture</i></p>
First name:	Eddy	
Family name:	LeDuca	
Gender:	Male	
Birthday:	01/01/1970	
Institution:	Open University of	
The Netherlands		
City and country:	The Hague,	
The Netherlands		
Contact information:	eddy.leduca@gmail.com	

About me *Tell what you want to tell about yourself.*

I started my studies in Environmental and Natural Science at the OUNL almost two years ago. Now, I am starting with the Master. I considered myself as a highly motivated person, who always looks at the positive side of problems. Problems are opportunities to learn!

Interests and hobbies *Tell what you want to tell about your non-work related areas of interest and hobbies.*

- Motor cycle restoration and motorcycle touring
- Tutoring web-based photography courses about picturing nature life.

Expectations of EES Community *Insert what you expect of EES.*

I am looking forward to EES as a new experience. I am looking forward to meet new people from different cultures working in the same domain and acquiring new insights on environmental policy issues.

Expertise areas *Tell in what fields you have expertise and how your peer EES members can contact you on these subjects.*

I can provide help based on my expertise in business administration and policy analysis. I also have knowledge in sustainable agriculture, eco-design (the basics) and energy crops.

Fields of interest *EES issues you are interested in. Fields where you have no or not so much expertise, but in which you are interested and want to learn more.*

I would like to learn about European policies for sustainable development.

Learn and work experiences *Shortly describe your relevant prior learn and work experiences.*

Currently, I am working as a policy analyst at TsA (Tomorrow-s-Aqua), which is an environmental consultancy firm, specialized in water management. I've a master degree in management science in the University of Rotterdam, and I am currently enrolled in the Master in Environmental Science of the Open University of The Netherlands. I also took some courses about biochemistry, although I never got an official diploma in this area.

Suggestions *Present ideas and links to websites that are of interest to group members.*

<http://www.sustainability.com> and <http://www.worldbank.org>

Fig. 3.1 pEXPi profile template, a worked-out example

with others. The functionality to share and distribute learning activities should support inputs from different sources such as, for instance, blogs, photos, maps, and devices as mobile phones, digital cameras, PDAs and the like. This includes on-the-spot contributions to the community, such as described by (De Jong et al. 2007), combining blog and mobile functionalities.

Sharing and distributing is but one step away from the next guideline, which marries sharing with actual interaction.

3.3.2 Facilitate Participants' Interaction with Others and Support Knowledge Co-construction Between Them

To facilitate interaction and reactions to contributions by others, participants should be able to comment on each other's resources and profiles, recommend a learning activities or a contact to someone else, or share a set of favourite learning activities. Eddy, for instance, has a collection of favourite links to electronic resources, articles he has written, figures, etc. which he considers relevant to environmental sustainability. He should be able easily to include them (or refer to them) in the online learning community, to share them with the other participants and recommend them to his colleagues. Moreover, Eddy needs mechanisms to find people, resources and communities and visualise and browse the relationship between them.

The next step up from informing is knowledge co-creation. To support this, participants should be able to develop a common ground first. This is not easy; to rethink one's own position and look for new arguments in the debate requires processing new ideas and adaptations of one's own ideas. Community members with interesting different perspectives on the topic under investigation can trigger fruitful discussions. However, divergent ideas sometimes lead to irritation and conflicts. Affordances and interventions are needed, therefore, to reduce miscommunication, trigger effective knowledge communication and facilitate shared understanding.

This kind of support can be provided using tools for discussion and knowledge representation such as *Knowledge Forum* (Scardamalia and Bereiter 2006), *Belvédère* (Suthers 2001) or *TC3* (Text Composer, Computer supported and Collaborative tool) (Kanselaar et al. 2003). Another example of such a tool is the *Ideasticker* (see Fig. 3.2). Developed at the Open Universiteit Nederland, the *Ideasticker* is a mediating tool that provides anchors for exchanging ideas, and ultimately, knowledge co-construction. It is a post-it like tool that can be activated any time by the learners in a community or can be programmed to automatically pop-up in a separate window. There are several other such tools available, for an overview see (Van Bruggen 2003). To explain what these tools are capable of, we delve a little deeper into *Ideasticker*.

Ideasticker provides a simple format to articulate the ideas (proposal field) and its underlying motivation (motivation field). The participant also specifies her position regarding the effects of her idea or proposal (position field), and can indicate what type of reaction she is looking for (expectation field asking, for example, elaboration of the idea, approval or rejection). The structuring options are available for posting

Fig. 3.2 Ideasticker screen

new ideas as well as for replying to existing ones. In this way it becomes immediately clear on which aspect of the idea the other participant holds different views, to which aspect she agrees or for which aspect she wants additional information.

Furthermore, a social chat option with ‘emoticons’ is included in the Ideasticker form. Emoticons are used to communicate emotional content in written or message form (<http://en.wikipedia.org/wiki/Emoticon>); they might help to express nonverbal cues beyond what is found in the verbal text of a message, and may enhance the exchange of emotional information (Derks et al. 2007), which facilitates participants to express their feelings and fosters interpersonal trust formation (see also Sect. 3.4.1 below).

Ideasticker is a tool that can be used across different learning activities. For example, to exchange ideas between participants about climate change risks and the pertinent European regulations; or to facilitate teams to work collaboratively solving ill-structured problems, such as the development of a research plan or the definition of recommendations to tackle an environmental issue. Imagine that Eddy and five more participants decide to analyse a case about the impact of global warming in the Arctic. The team wants to investigate the problem, develop a collective solution, and report their findings including policy recommendations at European level, and their

impact for stakeholders. Using Ideasticker this team can discuss ideas, solutions and policies in a structured way and avoid misunderstanding while collaborating.

Distributing materials to others and collaborating with others in knowledge co-construction activities, raises the issue of how to keep track of contributions, whether they are made by you or by others, either alone or in a collaborative effort. This issue, as well as the matter of the privacy arrangements to go with it, is addressed in the two guidelines to follow.

3.3.3 Help Participants to Classify and Evaluate Their Own Contributions but Also Those from Others

Sustainable online communities need tagging functionalities, so participants can classify, explore, and organise not only their own contributions – such as learning activities, contacts and communities – but also those from other participants. To make tagging even more powerful, rating functionalities, should be included, allowing participants to rate their own learning activities, contacts and communities, but also those from others. Tagging and rating functionalities are part of the *self-categorisation dimension*.

To give an example, by tagging the informal online community ‘EES courthouse’ with tags like ‘EU-law’, ‘sustainability’, ‘environmental’, and so on, EES participants would be able to find the community more easily. Likewise, EES participants could find Eddy using tags such as ‘policy-analysis’ or ‘energy-crops’, so he could be identified and found as knowledgeable in these topics. Others could rate Eddy’s contributions, once identified, as useful or helpful, but also as mistaken or confusing; or a numeric system could be used (1–5 points), which may be less likely to cause offense. Whichever, system is chosen, as much as labelling helps find resources, ranking them helps choosing between what has been found.

3.3.4 Allow Participants to Control the Level of Privacy of (Their) Contributions, as Well as to Decide Whether These Are Offensive or Not

The *self-regulation dimension* addresses participants’ control of existing resources and communities. It covers functionalities to control the level of privacy of learning activities and communities; but also functionalities that allow participants to flag that a particular contribution is seen as offensive. For instance, when Eddy creates learning activities, he should be able to define who else can modify them; if they are available to anyone; perhaps even if others can rate them. Moreover, in some instances it may be necessary to allow originators of online communities to define who can join the community, modify its characteristics, or add learning activities.

Regarding offensive behaviour, participants should have the possibility to condemn unacceptable learning activities or communities. Imagine that a participant of the 'EES courthouse' community includes a resource that applauds xenophobic views. This may constitute an offensive – the participant has put it there to circulate his own xenophobic opinions – or not – the participant uses it to exemplify a particular line of argument he does not necessarily endorse himself. Distinguishing between the two can be tricky. This is why a community preferably should deal with this kind of issues itself. However, as a last resort, it should always be possible to call the network administrators to the rescue and, for instance, have the offender banned from the Learning Network.

We have explained guidelines that foster sustainable online communities and related them to functionalities that should be present in online learning communities. In the next section, we will focus on how interpersonal trust formation processes amongst participants could be supported. This leads to the formulation of another four guidelines.

3.4 Guidelines to Support Interpersonal Trust Formation in Online Communities

Infrequent interactions, undesirable behaviour such as free-riding and sucker effects, personal and task-related conflicts and low-quality knowledge building (Häkkinen 2004), may all cause the absence of interpersonal trust in online collaboration (Furumo and Pearson 2006; Walther 2005). Here, interpersonal trust is defined as the positive expectation a participant has about the intentions and behaviour of another participant in the community (Mayer et al. 1995; Rousseau et al. 1998). Interpersonal trust is desirable as it improves willingness to interact, to share knowledge, and to participate in processes for knowledge co-construction, all covered in the previous section.

In face-to-face encounters, people form an impression of others based on the different types of signs and signals they acquire through their senses (e.g., a sound, a smell, a touch, a gaze) (Kandola 2006; Riegelsberger et al. 2004), either by first-hand experience or through other people (Hung et al. 2004). With the help of these signs and signals, they assess to what extent a person is reliable (Hardin 2002). This 'first' impression is not only grounded in these signs and signals gathered, but also coloured by one's existing cognitive schemata (e.g. stereotypes) and one's general trust propensity towards people (Bacharach and Gambetta 1997; Hung et al. 2004). For example, certain countenances are perceived to be more honest than others. This interpretation of information is part of everyone's implicit personality theory (Arnold et al. 1998).

In most online contexts, however, people have fewer and different types of signs and signals available. They lack information to form a first impression and to estimate the trustworthiness of other participants, on which they can start building their trust. This does not imply that it is impossible to form any level (either positive or negative) of interpersonal trust, it only means that it will be harder and take longer (Walther 1995, 2005).

Table 3.2 summarises the guidelines to support interpersonal trust formation. In the remainder of this section, we discuss them in turn.

Table 3.2 Guidelines to support interpersonal trust formation in online communities

Route through which information is acquired in face-to-face situations	Guideline
First encounter	Promote the exchange of off-task personal information
Indirect experience of others with a person (word of mouth)	Show and exchange information about participants' reputation
Direct personal experience with a person during collaboration	Show information about participants' presence, activities, and availability to the rest of the community
Contextual information	Show information about community's characteristics

3.4.1 Promote the Exchange of Off-Task Personal Information

In such mediated settings as an online community, some signs (e.g. body stature, odours) as well as opportunities to collect these signs (e.g. a coffee break between meetings) are missing. Without this information, people might use their existing schemes as a principal means to form impressions and might sooner form stereotypical judgments (Hung et al. 2004).

Consequently, communities should promote other means to form a first impression and discover familiar, shared aspects, between each other. Kivimäki et al. (1998) found that if community members' conversations passed the point of short and discontinuous phrasing, they all followed the pattern of asking personal information, such as about gender, age, city of residence, daily professional and private activities. People, in fact, have a need to collect this type of information. Indeed, Jarvenpaa and Leidner (1998) found that the exchange of personal information between people in an online group positively influenced formation of trust.

Exchange of personal information can be facilitated in different ways. One method is to provide a template for a personal profile and invite people to write about themselves. Although we are still investigating what personal information is especially important to form a first impression, a first pilot with the above mentioned pEXPI template gave promising results (Rusman et al. submitted). Another method is to facilitate informal talks between community members, by creating a separate space (coffee room) in which social communication is supported with, for example, chat and email (Kandola 2006; Zheng et al. 2001). Members of the EES, for example, have a forum named *Coffee Room*, in which they can interact informally. Regularly, one of the participants posts a comment regarding some out-of-context topic (e.g., music, cartoon, movie, etc.), and asks other participants to react and share their thoughts about it.

Yet another approach is to design off-task activities, called 'icebreakers' (Salmon 2003), in which participants introduce themselves to the rest of the community in an informal way (see also Cheak et al. 2006). Several formats to do this are available,

e.g. by asking participants to introduce their peer participants to the rest of the community, explaining their interests, posting a picture they would like to share, etc. This way, the presenter has to get in touch with the person she has to present to the rest of the community and discuss how this will be done (Berlanga et al. 2007a). For instance, imagine that Eddy LeDuca has to be presented to the rest of the community by other member, they might have to get in touch via e-mail or chat so Eddy can mention he would like to be presented as a motorcycle fan who has just bought a old Moto Guzzi V7, and maybe he would also like to share with the rest of the community a picture of his new motorcycle.

3.4.2 Show and Exchange Information on the Reputation of Participants

In face-to-face situations in which people do not know each other in advance, but have the opportunity to ask others about the characteristics and experience with the unfamiliar person in a context, they will certainly use this extra opportunity to gather information. The relation of the information gatherer with the information provider determines the weight and importance of the information gathered. Such *reputational information* is especially well spread and readily available within strongly knit communities: communities in which the closeness and interaction frequency is an indication of strong ties between members (Levin et al. 2004). In order to spread information, connections should exist between participants in a community and people should have the opportunity to build a reputation. This is not always the case. Reputations built within one context are not necessarily visible in other, context-similar communities. Communities then should provide ways to make connections between people visible, indicate the strength of their relations and provide space for recommendations of others, who becomes part of their 'online identity'. In order to make this work, an 'online identity' should be verifiable, stable and transferable from one community context to another (Kollock 1998; Riegelsberger et al. 2004). It is important to bear in mind, nonetheless, that privacy issues will arise when participants' online identity and reputation are transferred from one community to another.

Apart from the availability of reputational information, the fact that a participant is part of an existing social order will provide an incentive for more trustworthy behaviour (Riegelsberger et al. 2004). It will help people to behave 'nicer', i.e. be more inclined to collaborate. This not only works between two people, but also in a Learning Network. Between two people, the chance that a person misbehaves is reduced if there is a fair chance that they will meet in the future (and depend on each other again), due to a higher chance on reciprocity of favours (or mischief). In a Learning Network, this works more or less the same, but the behaviour will then depend on the chance that the 'bad' or 'good' reputation will travel within the network. For instance, take for instance the Master in Environmental Science that Eddy is taking. Participants of this Master know they will, for sure, interact until the

Master finishes. Afterwards, they are less likely to meet. Conversely, interactions in informal online communities, such as the EES, are not limited to a particular period; participants can never be sure that no interaction will occur in the future, so they might behave ‘nicer’.

3.4.3 Show Information About Presence, Activities, and Availability of Participants to the Community

In face-to-face situations, the level of interpersonal trust formed at the initial stages of a relationship between people will evolve through time, based on the direct positive or negative experiences a person has with others while collaborating. It can range from more cognitively-based trust to more affectively-based trust (Kanawat-tanachai and Yoo 2005). Also the stability of the image formed will improve: whereas people are prone to change this image after their initial encounters, after a while it will become a relatively stable construct, not easily changeable. This ‘final’ image will mainly be based on the behaviour of the other, e.g. reactions during possible task or personal conflicts, timely and high-quality contributions to a project, communicative behaviour during personal, difficult times.

Based on this information a participant ‘grows’ a personal history with the other and fills existing gaps in their initial image and impression. This will also happen in online collaborative situations, although due to different communication channels and lacking contextual clues, this information may be more sensitive to misinterpretation. People often do not ‘automatically’ know or realise when the other experiences a ‘busy’ period, is available at a specific moment to do work on the project, or is working in a different time zone, and can easily misinterpret ‘silence’ (Kandola 2006; Walther 1995, 2005).

In addition, sometimes, it can be hard to keep an overview of all communication and to attribute achievements to the right persons: who contributed to what and what appointments were made with whom and were they really met? To help prevent this kind of misconceptions, information on the online presence of participants, their general availability for an activity over time and their contributions should be made visible. Online presence information will also help to increase the likelihood of informal talks (e.g. by means of chat), which also has a positive effect on interpersonal trust formation (see also Sect. 3.4.1).

3.4.4 Show Information About the Characteristics of the Community

In face-to-face situations, people also derive information used for image formation of an individual from the characteristics of the group he or she belongs to. If, for example, a participant of the EES is an employee of the European Commission, people might attribute such properties as ‘able’, ‘honest’, or ‘well informed’, based

on their conceptualisation of the organisation at large. In order to do this, participants should have or be enabled to form an image of the groups characteristics (Riegelsberger et al. 2004). Therefore, information on the purpose and the overall community activity should be made visible upon entering a community, so people can judge right from the start what group membership of this group would imply: is it a community they could feel at home with because they have things in common with its inhabitants and the group is a lively one? Also, contributions that are highly valued within the community should be made visible, so that first-time visitors can judge the overall community quality and compare it with their own standards and expectations.

Lastly, general rules and policies of the community should be made visible so that in case of misbehaviour people know what to expect (Kollock 1998; Riegelsberger et al. 2004). A distinction should be made between rules and policies that pertain to the Learning Network as a whole and its individual communities. At the Learning Network level, rules apply as to refrain from xenophobic postings, discussed above. At the community level, rules and policies are particular to the community in question. The former rules need to be explicit; the latter will often stay implicit. In the latter case, it pays to make them explicit so that participants know what reprisals to face in case of malfunction. Offending rules at the network level may result in dismissal, offending rules at the community level typically has less severe consequences. Whatever the case, this type of consequences of mischievous behaviour will be an extra incentive for trustworthy behaviour within the group, but will also become part of the general trust model of a participant. It works as an extra guarantee.

3.5 Conclusion

In this chapter we have argued that interaction in online learning communities will not emerge automatically, so special affordances should be provided. We singled out community sustainability and interpersonal trust formation as important aspects to be considered in the set up of services for online learning communities.

We introduced guidelines to foster sustainable online communities. These claim that the Learning Network should offer participants functionalities in order to be able to (1) manage their own presence and contributions in communities, (2) to organise the community contributions and support knowledge co-construction, (3) to classify and evaluate participants' contributions, and (4) to regulate and control contributions. Furthermore, we also explained that in order to foster interpersonal trust formation in online communities, provisions have to be taken to counterbalance the lack of signals and signs normally perceived in face-to-face situations. To this end, a Learning Network should contain services that allow participants in online learning communities (5) to promote the exchange of off-task personal information, (6) to show and exchange information about participants' reputation, (7) to show information about participants' presence, activities and availability to the rest of the community, and (8) to show information about community's characteristics.

References

- Arnold, J. et al.: *Work psychology: Understanding human behaviour in the workplace* (Financial Times Professional Limited, Essex 1998) p 524
- Bacharach, M., Gambetta, D.: Trust in signs. In: *Trust in society* ed by Cook, K.S. (Russell Sage Foundation, New York 1997) pp 148–184
- Berlanga, A.J. et al.: Design of e-learning interaction spaces: Lessons learned from practice. In: *UniverSALearning 2007. International Conference on Technology, Training and Communication. Extended Papers*, Salamanca, Spain, September 12–14 (2007a)
- Berlanga, A.J. et al.: Functionality for learning networks: Lessons learned from social web applications. In: *ePortfolio*, (Maastricht, The Netherlands, October 18–20 2007b)
- Berlanga, A. et al.: Ad hoc transient communities: Towards fostering knowledge sharing in learning networks. *Int. J. Learn. Tech.* **3**(4), 443–458 (2008a)
- Berlanga, A.J. et al.: On the importance of personal profiles to enhance social interaction in learning networks. In: *Proceedings of Web Based Communities Conference* ed by Kommers, P. (IADIS Press, Amsterdam, The Netherlands, July 2008b)
- Bitter-Rijkema, M.E. et al.: Supporting knowledge elicitation for learning in virtual teams. *Educ. Tech. Soc.* **5**(2), 113–118 (2002)
- Bitter-Rijkema, M. et al.: Elicitation support requirements of multi-expertise teams. *J. Interact. Learn. Res.* **16**(2), 133–154 (2005)
- Boyd, D.: Friends, friendsters, and top 8: Writing community into being on social network sites. *First Monday* **11**(12), (2006)
- Brouns, F. et al.: Personal profiling to stimulate participation in learning networks. In: *Proceedings of the ePortfolio 2007 Conference* (Maastricht, the Netherlands, October 18–20 2007)
- Van Bruggen, J.: *Explorations in graphical argumentation. The use of external representations of argumentation in collaborative problem solving* Unpublished PhD thesis (Open University of the Netherlands, Heerlen, The Netherlands 2003)
- Cheak, A.M. et al.: Enhancing the social network dimension of lifelong competence development and management systems: A proposal of methods and tools. In: *Proceedings of international workshop in learning networks for lifelong competence development* ed by Koper, R., Stefanov, K. (TENCompetence, Sofia, Bulgaria 2006) pp 117–125
- Cörvers, R. et al.: Virtual seminars – or how to foster an international, multidisciplinary dialogue on sustainable development. In: *Crossing boundaries. Innovative learning for sustainable development in higher education* ed by de Kraker, J. et al. (Verlag für Akademische Schriften, Frankfurt am Mein 2007)
- De Jong, T. et al.: Contextblogger: Learning by blogging in the real world. In: *ePortfolio* (Maastricht, The Netherlands, October, 18–20 2007)
- Derks, D. et al.: Emoticons and social interaction on the internet: The importance of social context. *Comput. Human Behav.* **23**(1), 842–849 (2007)
- Furumo, K., Pearson, J.M.: An empirical investigation of how trust, cohesion and performance vary in virtual and face-to-face teams. In: *39th Hawaii International Conference on System Sciences* (Hawaii 2006)
- Girgensohn, A., Lee, A.: Making web sites be places for social interaction. In: *Proceedings of the 2002 ACM conference on computer supported cooperative work* (ACM Press, New Orleans, Louisiana, USA 2002) pp 136–145
- Häkkinen, P.: What makes learning and understanding in virtual teams so difficult? *Cyber. Psychol. Behav.* **7**(2), 201–206 (2004)
- Hardin, R.: *Trust and trustworthiness* (Russell Sage Foundation, New York 2002) p 234
- Hung, Y.C. et al.: Trust in virtual teams: Towards an integrative model of trust formation. In: *37th Hawaii International Conference on System Sciences* (Hawaii 2004)
- Jarvenpaa, S., Leidner, D.: Communication and trust in global virtual teams. *J. Comput. Mediat. Commun.* **3**(4), 791–815 (1998)
- Kanawattanachai, P., Yoo, Y.: Dynamic nature of trust in virtual teams. *Sprouts: Work. Pap. Inf. Environ. Syst. Organ.* **2**(2), 41–58 (2005)

- Kandola, P.: *The psychology of effective business communications in geographically dispersed teams* (Cisco systems, San José 2006)
- Kanselaar, G. et al.: Designing argumentation tools for collaborative learning. In: *Visualizing argumentation: Software tools for collaborative and educational sense-making* ed by Kirschner, P.A. et al. (Springer, London 2003) pp 51–73
- Kivimäki, A. et al.: Identity in virtual communities. *SIGGROUP Bull.* **19**(3), 29–33 (1998)
- Kollock, P.: Design principles for online communities. *PC Update* **15**(5), 58–60 (1998)
- Koper, R., Tattersall, C.: New directions for lifelong learning using network technologies. *Br. J. Educ. Technol.* **35**(6), 689–700 (2004)
- Kreijns, K.: *Sociable CSCL environments. Social affordances, sociability, and social presence* Unpublished PhD thesis (Open University of the Netherlands, Heerlen, The Netherlands 2004)
- Kreijns, K. et al.: Identifying the pitfalls for social interaction in computer-supported collaborative learning environments: A review of the research. *Comput. Human Behav.* **19**(3), 335–353 (2003)
- Levin, D.Z. et al.: The strength of weak ties you can trust: The mediating role of trust in effective knowledge transfer. *Manag. Sci.* **50**(11), 1477–1490 (2004)
- Mayer, R.C. et al.: An integrative model of interorganizational trust. *Acad. Manag. Rev.* **20**(3), 709–734 (1995)
- Ogg, H. et al.: *Handboek samenwerkend leren digitaal ondersteund* (Digitale Universiteit, Utrecht 2004)
- Riegelsberger, J. et al.: The mechanics of trust: A framework for research and design. *Int. J. Hum. Comput. Stud.* **62**(3), 381–422 (2004)
- Rogers, P., Lea, M.: Social presence in distributed group environments: The role of social identity. *Behav. Inform. Tech.* **24**(2), 151–158 (2005)
- Rourke, L. et al.: Assessing social presence in asynchronous, text-based computer conferencing. *J. Distance Educ.* **14**(2), 51–70 (1999)
- Rousseau, D.M. et al.: Not so different after all: A cross-discipline view of trust. *Acad. Manag. Rev.* **23**(3), 393–404 (1998)
- Rusman, E.: Provide personal identity information. E-len project design patterns http://www2.tisip.no/E-LEN/patterns_listing.php?sig=3&show=Explore+SIG+patterns (2004) Cited 28 May 2008
- Rusman, E. et al.: Theoretical framework for the design and development of a personal identity profile fostering interpersonal trust in virtual project teams. In: *6th workshop on social intelligence design* (Trento, Italy, July, 2–4 2007)
- Rutjens, M. et al.: *Handleiding voor de inrichting van een ontwerpomgeving voor informatica-opleidingen* (Digitale Universiteit, Utrecht 2003)
- Salmon, G.: *E-tivities: The key to active online learning* (Kogan, London 2003) p 223
- Scardamalia, M., Bereiter, C.: Knowledge building: Theory, pedagogy, and technology In: *Cambridge handbook of the learning sciences* ed by Sawyer, K. (Cambridge, New York 2006) pp 97–118
- Short, J.A. et al.: *The social psychology of telecommunications* (Wiley, New York 1976)
- Suthers, D.: Towards a systematic study of representational guidance for collaborative learning discourse. *J. Univers. Comput. Sci.* **7**(3), 254–277 (2001)
- Vogten, H. et al.: Using the personal competence manager as a complementary approach to IMS learning design authoring. *Interact. Learn. Environ.* **16**(1), 83–100 (2008)
- Walther, J.B.: Relational aspects of computer-mediated communication: Experimental observations over time. *Organ. Sci.* **6**(2), 186–203 (1995)
- Walther, J.B.: The rules of virtual groups. In: *38th Hawaii International Conference on System Sciences* (Hawaii 2005)
- Zheng, J. et al.: Trust without touch: Jump-start trust with social chat. In: *CHI '01 extended abstracts on Human factors in computing systems* (Seattle, Washington 2001)

their achievements and their appreciation of learning in a Learning Network setting. In general, individual success or failure on a learning activity depends on the extent to which learners perceive themselves as genuinely participating in a community (Wegerif et al. 1998).

In formal learning settings, autonomous learners are likely to make extensive demands on their teachers. After all, autonomous learners do not come in cohorts or classes, nor do they have uniform learning paths and goals that may be captured in preset curricula. This learner heterogeneity is bound to lead to a great variety of requests for support. Because of the lack of an available social structure, autonomous learners cannot rely on each other's help either, which tends further to increase the educator's workload. What little evidence is available seems to support these arguments (Romiszowski and Ravitz 1997 as cited in Fox and MacKeogh 2003). Rumble (2001, pp. 81, 82) quotes as much as a twofold load increase. Ideally, an online teacher should mainly facilitate student learning processes, while a teacher in a traditional setting should primarily select and share content (Beaudoin 1990; Salmon 2004). So in online learning, the teacher is to provide the students timely with feedback regarding their learning process rather than the subject matter or their learning products (Hardless and Nulden 1999). In practice, however, online teachers are responsible for both the learning process and the learning product. The latter entails activities such as (1) grading, (2) initiating, receiving and responding to messages, (3) collecting and marking assignments, and (4) maintaining and updating course content (Beaudoin 1990). De Vries, et al. (2005) note that teachers in online and blended learning environments find initiating, receiving and answering questions of students time-consuming. In other words, online teachers receive numerous content-related questions that need to be answered. Taken together this pile of responsibilities could easily overload the teacher. In online learning, it is therefore of the utmost importance to provide services that enhance a student's learning process and yet do not increase the work load of teachers (Fox and MacKeogh 2003).

Non-formal learning is as much intentional as is formal learning; however it does not rely on the usual infrastructures that are so characteristic of formal learning: schools, curricula, classes and cohorts. In Learning Networks – that are designed to sustain non-formal learning although bouts of formal learning may occur – learners will have no fewer requests for support than do learners in formal settings; for example, they will have the same kind of content-related questions as learners in formal settings. However, finding people that can help out will be harder precisely because of the lack of a formal learning infrastructure. Therefore, specific services have to be made available that offer them appropriate support.

This chapter will focus on the question of what rules and policies are conducive to the emergence of an adequate support infrastructure in the non-formal settings of a Learning Network. Once such an infrastructure is in place, a variety of ways again embedded in rules and policies may be used to guarantee its *continued existence*. This was the subject of the previous chapter. Here the focus is on the phase of the *emergence* of social interaction and the formation of an incipient social infrastructure. It is in this stage that the lone learner makes acquaintance with his or her online peers. More specifically, since learning through the exchange of knowledge is the

ultimate objective of any Learning Network, we will present guidelines for effective and efficient *knowledge dating* in Learning Networks; ultimately, this should lead to *knowledge sharing*. During our discussion, we will introduce the notion of ad-hoc transient communities, temporal online gatherings of people focused on a particular issue, which precisely because of their focus and transience help build a social infrastructure in a Learning Network. Ad-hoc transient communities may be seen as a knowledge dating and sharing service, offered in the context of a Learning Network. However, before going into this, we will start our story by briefly recapitulating a case begun in Chap. 2.

Remember how Eddy LeDuca, Jannie Barends, Bas Timmer and Jessica Zwart were all motorcycle enthusiasts with a particular craving for vintage motorcycles. Eddy is a policy analyst at the environmental consultancy firm *TsA*. Recently, he bought an old Moto Guzzi V7 from 1972, which he wants to renovate. Jannie is an early-retired expert on the Moto Guzzi V7, about which she owns a whole library of manuals detailing how to maintain, rebuild and repair motorcycles. Bas is a car mechanic who dreams of running his own garage in the near future. And Jessica works for the research and development department of Moto Guzzi at Mandello del Lario, Italy. Eddy, Jannie, Bas and Jessica are really all lifelong learners who – from their various own perspectives – want to expand their knowledge about motorcycles, in particular the Moto Guzzi V7 from 1972. Would there be a better a way to serve their interests than they do currently, by rather haphazardly surfing the Internet and, every so often, engaging in discussion fora?

4.2 Knowledge Dating in Learning Networks

Lifelong learning professionals and amateurs such as Eddy, Jannie, Bas, and Jessica will want to control their own learning activities. They build, for example, their own learning plans (e.g., the renovation plans of Eddy), produce their own reports on assignments (e.g., the business plan of Bas or the research plans of Jessica), and collect their own (scholarly) references and bookmarks (e.g., the library of Jannie). Because of this desire for autonomy, they will oppose any attempt to being put in cohorts or classes or having to submit to a curriculum, as would happen were they to participate in formal learning arrangements. An unfortunate side-effect of their autonomy, however, is that they may easily stay isolated, lacking a focal point for interaction that they can build on and develop. If such learners do not feel they belong to a larger community or network, they will not easily interact with peers; which hampers the emergence of a deeper understanding of what each of them wants to achieve. Lack of peer communication may ultimately negatively affect their success as a lifelong learner and even as a professional. Research shows that individual success on learning activities depends on the extent to which learners perceive themselves to be insiders of a network (Wegerif et al. 1998).

Our answer to this predicament is that these learners should first of all become a member of a Learning Network, for instance one that is devoted to vintage motor cycle aficionados. Second, a mechanism should be in place, which is part of the

services offered by the Learning Network, that forges social ties between them. From the description of the case, it appears that Eddy, Jannie, Bas and Jessica would make good conversation partners, even good peer learners. What is missing is a means of building on this potential. Ad-hoc transient communities, we claim, are the way to seed social interactions between them, which, ultimately, lead to knowledge sharing.

Ad-hoc transient communities by definition have a highly specific focus and lack permanence. They serve a particular goal and do so only temporarily. A case in point would be finding an answer to Eddy's question, discussed earlier, on how to clean and repair the carburettor for his Moto Guzzi. The ad-hoc transient community is only about this question, it ceases to exist once it is answered. However, since the participants in this community know each other, this contributes to the social space in the Learning Network. The participants may continue their engagement with each other, or they may not. This is up to them. Thus they maintain maximum autonomy and control. In the next sections we will explain how this could be done. If they decide to stay in touch, this may be the beginning of a community within the vintage motorcycles Learning Network fully devoted to the Moto Guzzi V7 from 1972 only. Over time, others will join in and a community arises with a more or less stable membership and a more or less fixed set of topics. What matters here is that this V7 community then has emerged from user initiated ad-hoc transient communities that discussed issues related to the Guzzi V7. We will discuss in more detail how in our view this should work and what measures need to be taken to make it work.

4.3 Knowledge Dating in Ad-Hoc Transient Communities

Effective knowledge sharing within ad-hoc transient communities is dependent on effective knowledge dating. That is, for effective knowledge sharing to occur group members should be carefully matched with each other because some 'dates' work and some do not. Eddy, for example, should ideally team up with someone who could tell him how to repair the Guzzi carburettor instead of someone who is in the midst of finding out herself. So, first of all, measures are needed to match the right people. This leads to a few guidelines. First, to assure lively knowledge sharing in a community, it should consist of a *heterogeneous* group of people, such as veterans (e.g. Jannie) and newbies (e.g. Eddy), or lurkers (e.g. Bas) and posters (e.g. Jessica) (Preece et al. 2004). Second, specific *roles* should be recognisable to the community members to avoid problems during the knowledge sharing process due to confusion about who is knowledgeable and who is not (Greenwood et al. 1989). Obviously, these roles could change between ad-hoc transient communities. Jannie is knowledgeable about technical issues, but Jessica would probably be better able to advice on business plans etc.

4.3.1 Heterogeneity

In lifelong and professional learning, participants in any given domain of knowledge have different levels of competence, varying from novices (e.g. Eddy) to top-experts

(e.g. Jannie), from practitioners (e.g. Bas) to researchers and developers (e.g. Jessica). Traditionally, in formal learning settings, the heterogeneity of learners has been reduced as far as possible to allow class-based teaching. This is done by providing clear entry requirements and using cohorts or groups that are considered homogeneous. In lifelong and professional learning, necessarily the door is opened to exploiting the heterogeneity of learners by putting together ad-hoc transient communities in which novices collaborate with more experienced people.

The prosperity of any community crucially depends on the characteristics of the people in it. First of all, people differ with regard to their experiences with working in an online community. Often learners are divided in veterans and newbies. Brown (2001) found that veterans showed good community behaviour. They were supporting and encouraging peers, sharing knowledge and experiences, reflecting on past learning, and sustaining friendships and/or acquaintanceships begun earlier. Newbies, however, depended much less on other group members and were reluctant to call for tutor help. They preferred a tight class structure with frequent interaction and helpful assessment from the tutor. It seems therefore wise to populate an ad-hoc transient community with both veterans and newbies. Because of their experience, veterans model good community behaviour to the newbies. Newbies can turn to veterans for support and encouragement instead of to the tutor. Although this helps to create an online community, veterans need an incentive to continue to interact with newbies. Veterans are inclined to do their 'duty' in the beginning but after a while tend to restrict their communication to veterans only, which hinders the expansion of the social structure in the Learning Network (Brown 2001). In our case, for Eddy, it would seem wise to team up with Jannie, the Guzzi-Godmother. She not only has a lot of knowledge about the Moto Guzzi V7, but she also knows how the Learning Network works. She can answer his question about how to repair and clean the Guzzi-part. For Jannie on the other hand it might be less interesting to team up with newbies as Eddy because their questions may not be very challenging for her.

Second, most people are trend-followers, but it is the trendsetter how makes the difference. Nichani (2001) describes three types of trendsetters, each of whom could significantly influence the thriving of any community: connectors, mavens and salesmen. *Connectors* form the 'social glue' of a community; they are very sociable and attentive and have a talent for making friends. *Mavens* are the information experts that have a talent for collecting information and are willing to tell others about it. *Salesmen* are persuaders, they have a tendency to reach out to the unconvinced and persuade them to join the community. The absence of trendsetters in a community, which then consists of trend-followers only, will negatively influence elementary features such as belonging, trust and social interaction (see also Chap. 3). So, an ad-hoc transient community populated with Jannie (connector), Jessica (maven) and Eddy (salesman) would work very well. Adding Jessica to the group in order to try and solve Eddy's question might be a win-win situation for all. Eddy finds an answer to his question; Jessica gathers information on Moto Guzzi parts that need to be redesigned and Jannie learns from Jessica about new developments for the Moto Guzzi.

Third and related to the issue of trendsetting, participants of online newsgroups differ in their inclination to either lurk or post in a community. A *lurker*, by

definition, belongs to a community but never posts in it. The percentage of lurkers in established communities is very variable (i.e., ranging from 0 to 99%; Preece et al. 2004). For example, lurkers appear to make up 45.5% of health support communities while the lurker population in software support communities could be as high as 82% (for an overview, see Preece et al. 2004). Reasons for not posting range from ‘didn’t need to post’, ‘needed to find out about the group’ ‘couldn’t make the software work’ ‘didn’t like the group’ to ‘had nothing to offer’ (Preece et al. 2004). Posters and lurkers are attracted to a community and join it for the same reasons. However, *posters* feel their needs are better met, perceive more benefit and feel a greater sense of membership than lurkers. Partly because posters do not regard lurkers as inferior members, lurking is not necessarily a problem in existing, active communities. Without a critical mass of posters, however, a community will never thrive (Preece et al. 2004). In our example, it may not be wise to involve a lurker such as Bas in an ad-hoc transient community set up to resolve the carburettor question. He would be of no use as his impact would be negligible. On the other hand, he might change his behaviour if personally invited. In the communities of practice Wenger describes, newcomers become ‘socialised’ through a process of peripheral participation which gradually evolves into full-fledged participation (Lave and Wenger 2002; Wenger 1999).

Finally, heterogeneity in levels of knowledge can have different effects on learning. Although for example, King and colleagues (1998) found that *peer-tutors* do not necessarily have to be more competent or more knowledgeable than their tutee counterparts, a study of Hinds et al. (2001) indicates that tutors equal in competence convey qualitatively different knowledge than more distant tutors. The near tutors – those who are similar to their tutees in knowledge level – use more concrete statements during their interactions with the tutee. In contrast, the distant tutors – those with a higher level of knowledge – convey more abstract and advanced concepts. Heterogeneity in level of knowledge between learners thus leads to a wide spectrum of knowledge shared in the community. To what extent the composition of an ad-hoc transient community needs to be heterogeneous in this respect depends on the goal of the community. It may well be that certain types of knowledge sharing or certain occasions or objectives maximally benefit from heterogeneity while others definitely do not. As discussed earlier, it may be sensible for veterans as Jannie and newbies as Eddy to team up, but there is a fair chance that in the long term the veterans do not find these interactions challenging any longer. In this case, it might be better for Eddy to team up with a knowledgeable peer, that is, someone whose level of knowledge is similar to his but yet knows the answer to his question.

4.3.2 Roles

Almost by definition, a Learning Network lacks the structure of the traditional class. Since no or few teachers are available to share and transmit knowledge, learners have to take up this role themselves. This change of role between learner and teacher could generate a lot of ‘noise’ in a Learning Network. That is, it could cost Learning

Network participants much effort without contributing to and even interfering with the knowledge sharing process. In our case, for example, imagine Eddy in his garage with his Guzzi V7 parts spread out over the floor, trying to find out how to reassemble the fuel system. Finally, he turns to his computer to visit an ordinary forum on vintage motorcycles to look for an answer. In it, Eddy has to go through threads that touch upon his problem. Failing to find an answer, he has to post his question in the hope that somebody answers it. Or he could in some other way try to find someone who is able to answer his question and directly approach this person. These processes are referred to as noise – or transaction costs, depending on your disciplinary background – since they do not themselves contribute to the knowledge sharing process; the more effort they cost, the less effort can be spent on the knowledge sharing process itself. Hence, in Learning Networks measures need to be taken, in the form of appropriate services that maximally reduce the noise or transaction costs so that the cooperative process of knowledge sharing can take off.

Improvements can be made by ‘filtering the noise’. That is, learners should not have to invest effort in activities that are not relevant for knowledge sharing. Learners who carry out complex cognitive task – such as for the first time in your life reassembling the fuel system of a Guzzi V7 – are expected to profit from a filtered network. These learners, as compared to learners that solve less cognitively demanding tasks, need all their cognitive capacity to solve the problem at hand, so, all effort that has to be invested in activities not relevant for problem solving is wasted and will hamper learning. It is surmised that providing Learning Network members with an identity based on their experience in the domain (i.e. novices and experts) or in the Learning Network (i.e. veterans and newbies) could enable them to take up roles. As a consequence, a structure for knowledge sharing could arise. For instance, domain experts can take up the role of tutor and domain novices will adopt the role of tutee while they all use a peer-tutoring format for knowledge sharing. In our example, if the Learning Network holds data about its members on their expertise and experience, it could bring suitable people together given a certain goal (the exact way in which this service operates will be discussed in Chap. 5). In this way, Jannie, Jessica or Bas could have joined Eddy in an ad-hoc transient community addressing his question. This would save him the trouble of searching through forum threads, seeking for someone able to answer his question, etc.

Greenwood et al. (1989), who carried out a longitudinal study on peer-tutoring in a formal learning setting (i.e. a classroom), found that students in peer-tutoring classes were more engaged in learning activities and knowledge sharing and less engaged in structuring these activities. Therefore, the students who learned within a peer-tutoring structure achieved higher learning outcomes than students who learned without it. Interestingly and importantly, the peer-tutoring structure seems to help both the peer-tutor and the tutee to achieve higher learning outcomes (Gyanani and Pahuja 1995). Effective peer-tutoring structures provide the tutors with support (e.g. probing/review questions, hints) that helps them effectively to guide the tutee’s learning process. King et al. (1998) state that such tutor-support raises the knowledge sharing between tutor and tutee to a high cognitive level that includes mutual exchange of ideas, explanations, justifications, speculations, inferences, hypotheses,

and conclusions. *Mutatis mutandis*, such effects should also pertain to tutor-tutee relations in ad-hoc transient communities, but certainly in the communities that emerge out of them.

4.4 Knowledge Sharing Through Ad-Hoc Transient Communities

Knowledge dating is but a stepping stone towards knowledge sharing, which is what Learning Networks are about. Ad-hoc transient communities, again, are the prime mechanism we propose to set knowledge sharing in motion. What conditions should be met for knowledge sharing in ad-hoc transient communities to occur? A survey of the literature (see Kester et al. 2007) yields two important conditions. First, to achieve and maintain social interaction, one should establish their recognisability, a historical record of their actions, and continuity of contact (Kollock 1998). This is the *accountability* condition. Second, ad-hoc transient communities should have a clear *goal* for knowledge sharing to occur.

4.4.1 Accountability

A sound social space is characterised by affective work relationships, strong group cohesiveness, trust (i.e., perceived reliability of the word of other group members and genuine interest in the welfare of group members), respect, belonging (i.e., recognition of membership) and satisfaction (Kreijns 2004; Nichani 2001; Rovai 2002). Social interaction enhances the emergence of social space. Task-driven interaction directed towards the completion of assigned tasks, however, could negatively influence the growth of this social space. When, for example, as part of the goal of a particular ad-hoc transient community, it is the members' task to assess each other, fear of criticism or reluctance to criticise could interfere with feelings of trust (Rovai 2002). Furthermore, mistaken expectations of what a community should bring, could also negatively influence social interaction and hence the emergence of social space. According to Brown (2001), individuals who felt that people needed to join voluntarily or felt that face-to-face association was necessary, only developed a sense of belonging and trust if they joined a face-to-face community of their own volition. So social interaction and, as a consequence, the emergence of social space is facilitated when socially and emotionally driven interaction is stimulated instead of task-driven interaction; the same facilitation is observed when people's expectations about a community are fulfilled.

More generally still, three social prerequisites should be met in order for social interaction, in particular cooperation, to occur: (1) all individuals must be able to identify each other (recognisability), (2) all individuals must be able to know how any other person has behaved in the past (history), and (3) any two individuals must be likely to meet again in the future (continuity). If individuals only meet once, they

are very much tempted to behave selfishly, which negatively influences the cooperation process. In addition, if individuals are not identifiable and no history of a person's behaviour is available, group members are more likely to act selfishly because they cannot be held accountable for their actions (Kollock 1998) (see Chap. 2 for a different tack on these desiderata).

First, the recognisability of learners can be ensured by forbidding the use of (multiple) aliases such as screen names. However, this may be hard to accomplish in practice. So, if one does not want to ban pseudonymity entirely, learners that go by a pseudonym should adopt a persistent one. Persistence can be guaranteed by using the learner profile as proxy for someone's identity (see Chap. 3). Note that different personal details can be made visible to different actors, depending on, for example, the privacy of the learner, the goal of the community, or the role of each learner in the community. So in the vintage motorcycles Learning Network, Jannie adopted the pseudonym 'Guzzi godmother' when she signed up for it. This pseudonym does not directly tell newbies that she is an expert when it comes to the Moto Guzzi V7 and a veteran in the Learning Network. However, Jannie's privacy settings allow every visitor of the Learning Network to take a look at a general part of her personal profile in which she flags her expertise and experience of the Moto Guzzi and her long involvement in the network. Her personal profile may also contain a lot of personal photos. Jannie's privacy settings allow specific members of the network to access these.

A historical record of user activities can be maintained by logging all the learner's activities. The ones significant for knowledge sharing, for example those that reflect content competence or knowledge-sharing competence, become part of the learner's profile. Content competence reflects the learner's mastery of the content within the Learning Network. To underpin content competence, the profile contains the products that resulted from a learner's activities, such as papers, reports, or even assessments. Knowledge-sharing competence refers to the ability of a learner satisfactorily to support peers during a process of knowledge sharing. This information could be acquired directly, by letting learners rate each other's performance in the ad-hoc transient communities they both participated in; or indirectly, by monitoring the achievements of the learners with whom the learner to be rated has shared his or her knowledge. Learner profiles should also incorporate this information. Furthermore, the parts of the personal profile that are accessible to all network members, should display the ad-hoc transient communities these members were active in.

Continuity of contact within a Learning Network can be guaranteed by the ad-hoc transient community structure that is implemented. Furthermore, the ad-hoc transient communities will continuously surface in the Learning Network to serve different purposes and, although they continuously change with regard to their composition, learners will likely meet again some time in some newly started ad-hoc transient community or other. Because Eddy had this problem with repairing the carburettor of his Guzzi V7 he came in contact with Jannie and Jessica, who appeared to be the right persons to help him given their personal profiles. Since especially Jannie has a lot of expertise with motorcycles, it is likely that Eddy and she will meet again when Eddy runs into another problem during his renovation project.

4.4.2 Goal Orientation

Ad-hoc transient communities should have a clear goal optimally to promote knowledge sharing. The goal could result from a request of a learner, for example, a content-related question (see Chap. 5). The learner request then forms the incentive for the process of knowledge sharing. Indirectly this request strongly influences the social interaction. One can imagine, for example, that a knowledge-sharing goal tied to answering content-related questions elicits interaction patterns different than a request to comment on a paper. Answering a content related question is a well-structured problem and therefore yields a limited amount of interaction. So if Eddy were to ask a question, Jannie and Jessica would give an initial answer, sometimes to be followed up by a few more exchanges of messages, to clarify the answer or even the question itself. Commenting on a paper, however, is an ill-structured problem and is likely to yield more and more prolonged interaction. Since there is no right or wrong answer to, for example, the question of what style or structure is best given the circumstances, answering this kind of request will lead to a fair amount of discussion between group members. So, the personal profile of Bas contains his business plan for his garage and upon his request an ad-hoc transient community arises to discuss this business plan. Jannie and Jessica join in and since Bas, Jannie and Jessica are all experts and thus know their business, it is likely that a vivid discussion will evolve around his request for help.

The social interaction pattern elicited by a request for knowledge sharing may be mediated by different interaction structures, many of which also characterise formal learning communities. Examples of such structures are *peer tutoring*, *Group Investigation* (Sharan and Sharan 1992), *Student Teams Achievement Division* (Slavin 1995), *Jigsaw* (Aronson and Thibodeau 1992; Bielaczycs 2001), *Structural Approach* (Kagan 1994) (each structure is a scenario to teach specific skills and, though not similarly articulated, it is implicitly assumed that situations are not identical), *Progressive Inquiry* (Rahikainen et al. 2001), the use of *scripts* (O'Donnell 1999; Weinberger et al. 2001), *scenarios* that prescribe collaboration activity (Wessner et al. 1999), *feedback rules* or requirements of a minimum degree of contributions to a discussion (Harasim 1993; Harasim et al. 1995). See Strijbos (2004, p. 33) for a detailed discussion. Strongly structured interactions, such as peer-tutoring or Jigsaw, are most suitable for knowledge-sharing goals that inherently elicit little interaction, because such interactions guarantee at least a minimum amount of social interaction. King et al. (1998) advocate a three-step structure for peer tutoring that consists of communication guidelines (i.e. listening, encouraging and giving feedback), an explanation procedure (i.e. the TEL WHY-procedure; telling in one's own words, explaining why and how, and linking of content), and questioning guidelines (e.g. asking comprehension questions or thinking questions). Jigsaw (Aronson et al. 1978; Slavin 1995) is a technique suited to situations in which students have to learn from written materials (e.g. textbooks, fact sheets). In Jigsaw, the academic content is divided up into as many sections as there are team members. They then have to study the section of the content allotted to them with members of the other teams, who are invited to study the same section. Together they form an 'expert group'.

After they all have become ‘experts’ on that section, they each return to their original team to share with it what they have learned (Kreijns 2004, p. 40).

Suppose Eddy, Bas, Jannie and Jessica are classmates instead of learners in a Learning Network. Eddy still has a problem with repairing the Moto Guzzi V7 carburettor and Jessica is assigned to help him. So, Eddy asks Jessica how to repair the carburettor and Jessica gives him the answer. In this example the interaction is very limited. The interaction could be increased for instance, by using the peer tutoring strategy of King et al. (1998). In this case, Jessica asks Eddy to explain the problem while giving him feedback, then they discuss the origin of the problem and then Jessica helps Eddy to find the solution of his problem himself by asking him relevant questions. Such a method increases the interaction between group members, which could enhance the effectiveness of knowledge sharing in terms of learning.

Weakly structured interactions, such as Progressive Inquiry, however, seem to be most suitable for knowledge sharing goals that inherently elicit much interaction. Progressive inquiry seeks to stimulate the same kind of productive practices of working with knowledge that characterise scientific research communities. By imitating the practices of these kinds of communities, students are encouraged to engage in extended processes of question- and explanation-driven inquiry. Accordingly, an important aspect of progressive inquiry is to guide students in setting up their own research questions and working theories. In practice, this means that students are making their conceptions public and working together to improve shared ideas and explanations. It is an essential element of this approach to constrain emerging ideas by searching for new information. Participation in progressive inquiry is usually embedded in computer-supported collaborative learning environments that provide sophisticated tools for supporting the inquiry process as well as sharing of knowledge and expertise. Imagine Bas, Jannie and Jessica are entrepreneur trainees and they are working on a business plan together. They write this business plan for the Chamber of Commerce where they have to register their business. Such a task already elicits a lot of discussion and the group would be helped if they were supported in explaining their ideas to the others, develop shared ideas et cetera by tools such as for example templates or external representations.

4.5 Conclusion

This chapter described how ad-hoc transient communities could be deployed effectively and efficiently to date for knowledge sharing. We discussed various guidelines for setting up such communities, the specific nature of the guidelines depending on whether they were about knowledge dating or knowledge sharing. Together, they form a recipe for increasing the success of ad-hoc transient communities. If successful, they should increase social cohesion of a Learning Network and give rise to the emergence of more permanent communities in it. Below, we summarise the guidelines discussed.

1. To foster knowledge dating, ad-hoc transient communities should be populated by a *heterogeneous* group of people (i.e. veterans and newbies, lurkers and posters or domain novices and experts).
2. To foster knowledge dating, the members of ad-hoc transient communities should have distinct and recognisable *roles*, for example, those of novice or expert.
3. To foster knowledge sharing, members of ad-hoc transient communities should be *accountable* for their actions. Therefore they have to be *recognisable* with a persistent identity; their *history* in the Learning Network should be made public and *continuity* of contact between members should be stimulated.
4. To foster knowledge sharing, each ad-hoc transient community should have a single, clear-cut goal, which may come in different kinds (e.g. content related questions or requests for comment).

References

- Aronson, E. et al.: *The Jigsaw classroom* (Sage, Beverly Hills, CA 1978)
- Aronson, E., Thibodeau, R.: The Jigsaw classroom: A cooperative strategy for an educational psychology course. In: *Cultural diversity and the schools* ed by Lynch, J. et al. (Palmer, Washington 1992) pp 231–256
- Beaudoin, M.: The instructor's changing role in distance education. *Am. J. Distance Educ.* **4**(2), 35–43 (1990)
- Bielaczycs, K.: Designing social infrastructure: The challenge of building computer-supported learning communities. In: *European perspectives on computer-supported collaborative learning. Proceedings of the 1st European conference on computer-supported collaborative learning* ed by Dillenbourg, P. et al. (Maastricht University, Maastricht 2001) pp 106–114
- Brown, R.E.: The process of community-building in distance learning classes. *J. Async. Learn. Network.* **5**(2), 18–35 (2001)
- De Vries, F.J. et al.: Identification of critical time-consuming student support activities in e-learning. *ALT-J Res. Learn. Technol.* **13**(3), 219–229 (2005)
- Fox, S., MacKeogh, K.: Can elearning promote higher-order learning without tutor overload? *Open Learn. J. Open Distance Learn.* **18**(2), 121–134 (2003)
- Greenwood, C.R. et al.: Longitudinal effects of classwide peer tutoring. *J. Educ. Psychol.* **81**(3), 371–383 (1989)
- Gyanani, T.C., Pahuja, P.: Effects of peer-tutoring on abilities and achievement. *Contemp. Educ. Psychol.* **20**, 469–475 (1995)
- Harasim, L.: Collaborating in cyberspace: Using computer conferences as a group learning environment. *Interact. Learn. Environ.* **3**, 119–130 (1993)
- Harasim, L. et al.: *Learning networks: A field guide to teaching and learning online* (MIT Press, Cambridge 1995)
- Hardless, C., Nulden, U.: Visualizing learning activities to support tutors. In: *CHI '99 extended abstracts on Human factors in computing systems* (1999)
- Hinds, P.J. et al.: Bothered by abstraction: The effect of expertise on knowledge transfer and subsequent novice performance. *J. Appl. Psychol.* **86**, 1232–1243 (2001)
- Kagan, S.: *Cooperative learning* (Kagan Cooperative Learning, San Juan Capistrano 1994)
- Kester, L. et al.: Facilitating community building in learning networks through peer-tutoring in ad hoc transient communities. *Int. J. Web Based Communities* **3**(2), 198–205 (2007)
- King, A. et al.: Mutual peer tutoring: Effects of structuring tutorial interaction to scaffold peer learning. *J. Educ. Psychol.* **90**, 134–152 (1998)

- Kollock, P.: Design principles for online communities. *PC Update* **15**(5), 58–60 (1998)
- Kreijns, K.: *Sociable CSCL environments. Social affordances, sociability, and social presence* Unpublished PhD thesis (Open University of the Netherlands, Heerlen, The Netherlands 2004)
- Lave, J., Wenger, E.: Legitimate peripheral participation in communities of practice. In: *Supporting lifelong learning* ed by Harrison, R. et al. (Routledge, London 2002)
- Nichani, M.: Communities of practice at the core. *Elearningpost* (2001)
- O'Donnell, A.M.: Structuring dyadic interaction through scripted cooperation. In: *Cognitive perspectives on peer learning* ed by O'Donnell, A.M., King, A. (Lawrence Erlbaum, Mahwah, NJ 1999) pp 179–196
- Preece, J. et al.: The top five reasons for lurking: Improving community experience for everyone. *Comput. Human Behav.* **20**, 201–223 (2004)
- Rahikainen, M. et al.: Progressive inquiry in CSILE environment: Teacher guidance and students engagement. In *European perspectives on computer-supported collaborative learning. Proceedings of the 1st European conference on computer-supported collaborative learning* (Maastricht 2001)
- Rovai, A.P.: Building a sense of community at a distance. *Int. Rev. Res. Open Distance Learn.* **3**(1) (2002)
- Rumble, G.: The costs and costing of networked learning. *J. Async. Learn. Network.* **5**, 75–96 (2001)
- Salmon, G.: *E-moderating: The key to teaching and learning online* (Taylor & Francis, London 2004)
- Sharan, Y., Sharan, S.: *Expanding cooperative learning through group investigation* (Teachers College Press, New York 1992)
- Slavin, R.E.: When does cooperative learning increase student achievement? *Psychol. Bull.* **94**, 429–445 (1995)
- Squires, D.: Educational software and learning: Subversive use and volatile design. In: *Thirty-Second Annual Hawaii International Conference on System Sciences-Volume 1* (1999)
- Strijbos, J.W.: *The effect of roles on computer-supported collaborative learning* Unpublished master's thesis (Open University of the Netherlands, Heerlen, The Netherlands 2004)
- Wegerif, R. et al.: Software design to support discussion in the primary curriculum. *J. Comput. Assist. Learn.* **14**, 199–211 (1998)
- Weinberger, A. et al.: Scripts and scaffolds in problem-based CSCL environments: Fostering participation and transfer. In *Computer-mediated cooperative learning. Symposium conducted at the 9th European Conference of the European Association for research on Learning and Instruction* (2001)
- Wenger, E.: *Communities of practice: Learning, meaning, and identity* (Cambridge, Cambridge 1999) p 336
- Wessner, M. et al.: Using learning protocols to structure computer-supported cooperative learning. In: *Proceedings of the ED-MEDIA '99 – World Conference on Educational Multimedia, Hypermedia & Telecommunications* (Seattle 1999)

brings suitable people together for a specific purpose, because it will allow the participant to focus on the knowledge sharing process by reducing the effort or costs. In the current chapter, we describe an example of a peer support Learning Network Service based on the mechanism of peer tutoring in ad-hoc transient communities.

Learners invariable have a lot of questions. These questions are of varying natures; some concern administrative or organisational matters, such as how do I enrol in this course, when is the exam; others are domain-related, such as what topic should I study next, or are content-related and arise when learners are studying a particular topic. In more traditional educational settings, there are provisions to cater for these questions. Administrative and organisational questions can be dealt with by the student administration services. Other types of questions are dealt with by the teachers, and tutors (De Vries et al. 2005). Turning to the teacher when a learner has a question, or asking a fellow student is more obvious in traditional settings, because in these circumstances learners usually know who partakes in the course. And there often are more opportunities, as there usually is a lot of face-to-face contact. In Learning Networks as we define them, it is more cumbersome for learners to obtain an answer to their question. The traditional support structures are lacking. In these circumstances, it is even highly likely that tutors are unavailable for several subject domains, e.g. because it is not economically viable, or because the resource provider did not provide for them. For certain questions people can be directed towards (online) resources, search engines or FAQs, but not all questions can be resolved in this manner. Some questions, such as content related questions, or requests for comments on a paper, or discussions on a certain topic can be dealt with only by contacting somebody else. Often people do not know whom to contact for a particular request.

This is where ad-hoc transient communities come in. In this circumstance an ad-hoc transient community will be set up to assist a learner who has a particular question. The participants of the ad-hoc transient community will consist of peers with the required expertise to support the learner in finding an answer to a specific question. The peers are selected by the peer support Learning Network Service, based on certain selection criteria that suit the goal and adhere to the guidelines for heterogeneity, roles and goal as described in Chap. 4.

Let's illustrate this with the case introduced in Chap. 2 Bas Timmer is one of the members of the vintage motorcycles Network. He currently works as a mechanic, but really would like to run his own garage. He already started with a business plan and already engaged in several lively discussions about the plan with his fellow Network members Jannie and Jessica. These discussions made him realise that technical skills alone are not sufficient to run a garage. He needs to learn how to manage a business and update his accounting and customer skills. Bas decides to join another Learning Network, on business administration, and starts with accounting. One of the first things to learn is how to make up a balance sheet. His maths is not very strong and he cannot figure out how to do the credits and debits properly. In particular, the assets accounts are confusing him. Just reading about it is not taking him any further and he decides to ask somebody for clarification. Because he effectively is on his own, he decides to call in the help of other people in the Learning

Network. However, he does not know whom to contact and turns to the peer support service of the Learning Network. Bas accesses the support request form. The request form guides him in detailing his question. The Learning Network Service creates an ad-hoc transient community, facilitated with relevant communication and collaboration tools, to which Bas and relevant peers with complementary expertise get access, with the goal to reach jointly at a satisfactory answer for Bas' problem.

5.2 Addressing the Solution

Our first attempt at implementing a peer support Learning Network Service focuses on content-related questions. The main assumption is that for the ad-hoc transient community to be effective we have to select suitable and eligible peers. The data on which to base this selection should be taken from a participant's profile. This requires that the relevant data is recorded into the profile, and that the Learning Network stores a record of every participant's progress through the Learning Network preferably into the participant's profile. Chapters 3 and 4 already indicate some of the data that should be stored. Below we explicate the data actually used in this implementation. In addition, the Service requires access to all documents that constitute the resources in the Learning Network.

5.2.1 Steps Involved

The model contains the following steps (Van Rosmalen 2008; Van Rosmalen et al. 2006, 2008b) as illustrated in Fig. 5.1.

Formulating a support request. A question form is available to the participant who wants to pose a question (tutee). This form contains guidelines on how to phrase the question in such manner that it provides maximum input to the peer tutors. Addi-

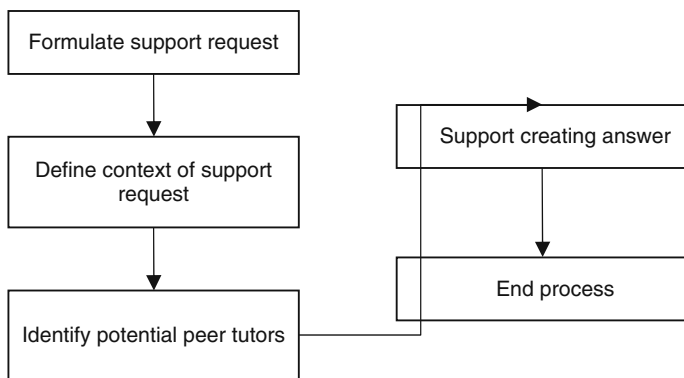


Fig. 5.1 The five steps involved in peer tutoring ad-hoc transient communities

tional information related to the question, such as urgency, date by which an answer is expected or required, subject domain of the question (if known) or the learning activity that raised the question, should also be provided on the form. Information related to the identity of the tutee can be determined automatically when the participants has to authenticate for the Learning Network or the form, else the form should request this data as well.

Defining the context of the support request. Usually somebody will have a question while studying a learning activity and the question will be related to that learning activity, but this does not have to be so. When somebody is studying more learning activities at the same time, in particular when these activities are closely related, it is not always possible to exactly pinpoint the activity that gave rise to the question. Sometimes a question relates to more than one learning activity. Or the question might be related somehow to the subject domain, but not directly to content in one of the learning activities. An indication of the context can be taken from the form. When a question form is available from within a certain learning activity, this can be automatically set on the form, but the tutee should be able to override this. When the form is also available outside the context of the learning activity, e.g. as a generic service at the Learning Network level, the form should contain a field in which the learner indicates the subject domain or learning activity to which the question pertains. Instead of forcing the learner to indicate always a learning activity, the option 'Do not know' can be used when the learner does not know exactly which learning activity the question relates to. The context of the question is rather important because it provides the basis on which suitable peers are selected. Therefore, we not only rely on the information provided by the learner but also have a mechanism in place to determine the learning activities that are relevant to the question. We use language technologies to assist in analysing the questions, which after all are formulated in natural language. Latent Semantic Analysis (LSA) is quite often used in an educational context (Van Bruggen et al. 2004; Landauer et al. 1998). LSA is used to find documents that are related because concepts are used instead of searching for the exact keyword. This allows people to use more of natural language in search queries instead of having to specify the correct keywords. We apply LSA to determine correlations between the text of the question and the documents in the Learning Network. Because every document belongs to a single learning activity, we can extrapolate to the learning activities the question is related to.

Identify potential peer tutors. This step involves finding and selecting suitable peer tutors and deciding upon the optimal number of peer tutors. The peer tutors are selected based on four criteria: content competence, tutor competence, availability and eligibility, as explained below. Most of the data required for the selection algorithms, such as activities taken, competences achieved, and learning goal, are already stored in Learning Networks and can be taken from a learner's profile, as well as system's logging data. Other data might require additional services that are not always present in a Learning Network; availability, for example, can be retrieved from a calendar, rating services need to be added to allow members to rate tutor's ability and quality of contributions.

Support for creating the answer. The previous step returns a ranked list of suitable peers. A number of these peers are then invited to assist in answering the question and are asked to join a collaborative workspace. The invitation contains the question and guidelines. The invitation or notification of the invitation preferably is sent out via e-mail, to ensure a quick response. The ad-hoc transient community is supplied with the necessary communication and collaboration facilities (such as a wiki or discussion forum), that allow the peers to jointly discuss and phrase an answer to the question. The question is made available, together with guidelines on how to proceed. To lower the peers' required effort, the system retrieves relevant text fragments related to the question from the resources in the Learning Network and makes those available in the ad-hoc transient community. These documents should assist the peers in getting a quick overview of documents relevant to the questions, and could also contain part of the answer or be the basis for the formulation of the answer.

End of process. After some time the peers should be able to formulate the answer to the question. The tutee decides when the process is ended; hopefully when the answer is satisfactorily. There may be several other reasons for a tutee's decision to end the process: because the peer tutors fail to provide an answer the tutee can agree with, because the process did not get started at all, or maybe just because it looks like no decent answer can be found in the near future, or they run out of time. Whatever the reason, it is up to the tutee to end the process. When the process is ended, the tutee should rate the peers and the product. The ad-hoc transient community then is dismantled and ceases to exist.

5.2.2 Peer Selection Criteria

We identified four criteria for the selection of peers who can act as potential tutors when it comes to answering content related questions. These criteria are designed to ensure quality and economy. Only suitable or competent peers should be selected, otherwise it will not be possible to get an acceptable answer to the request or query. In addition, peers should be able to meet the request and there should be a fair spread in workload. The weighted sum of these four criteria results in a rank-ordered list of suitable tutors. More details on the algorithms can be found in Van Rosmalen et al. (2008b).

- *Content competence.* First, the peer tutor should be sufficiently competent and knowledgeable in the subject domain in order to answer the question.
- *Tutor competence.* The peer tutor should have sufficient tutoring capabilities to be able to support the tutee.
- *Tutor eligibility.* The peer tutor should be eligible to answer the question.
- *Availability.* And of course, the peer tutor should be available to respond to the request within a certain timeframe.

Content competence is a measure of how well a peer has mastered the content relevant to the question. This information can be obtained from the member's profile in which evidence for competences attained is stored or in which the learning activities of the member together with completion status and competence level are recorded. However, competence assessment is not trivial, and even when the profile contains some information regarding competence, the interpretation of these data is not straightforward, let alone determining the competence proficiency level the learner obtained. An alternative would be to consider completion status of learning activities. Usually learning activities are designed for a certain learning objective or competence, and an assessment is used to ascertain whether the learner has met the objective. The result of the assessment can be taken as an indication of proficiency level. In our model, the content competence is expressed as a weighted sum of completion status of all learning activities that are relevant to the question. For a more detailed measure, time since completion can be taken into account.

Tutor competence indicates the ability of a peer to act as a tutor and is a measure of tutor capabilities. It refers to the capability to satisfactorily support the learner who has questions in such way that the tutee is satisfied with the process and the product. By nature, the staff tutor would be the most suitable tutor. Similarly, we can assume that learners who have progressed furthest in the Learning Network have obtained the most content competence and therefore are likely to be the most suitable peer to select. In addition, the peers that have been selected before or were more often selected, are likely to be better tutors because they gained experience in previous situations. Tutoring competence is difficult to measure because the evidence is hard to record. Another way to obtain a measure of tutoring competence is with a rating mechanism, in which the members rate both the tutoring process and the final answer. Or it could be derived from the frequency, size and quality of the contributions.

Tutor eligibility is a measure that is required to ensure a proper spread of workload amongst the members of the Learning Network. To avoid that the same people are selected time and again, the tutee should be compared to the potential peer tutors and only those peers that are nearest to the tutee should be selected. The nearest peers are those members that have a similar learning path, or at least have progressed as far. This can be measured by looking at the goal, i.e. the learning activities a member needs to complete to obtain the chosen competence and the completion status of the learning activities in the learner's profile. Peers who have completed the same learning activities (or get as close as possible to the tutee's profile) are most eligible.

Availability refers to the actual availability of the peers. Is the member present and active in the Learning Network, taking into account absence for holidays, days off, illness, etc., and workload in the sense of study-load, i.e. studying for exams or preparing an assignment, and previous participation (acting as tutor)? This measure is also time-dependent: more recent workload should have a higher weight or affect availability more than workload in the past.

For other types of support requests, these four criteria equally apply. The underlying algorithms might vary or be based on different parameters. In some cases the weight of the four criteria might need to be adjusted. Content competence is much less relevant when answering questions related to administrative matters. And when asked to answer a question about time and place of the assessment content and tutoring competence hardly are required. Nevertheless, the same principles apply and the selection criteria are still involved.

5.2.3 A Tutor Locator: Example Implementation of the Model

The model we described above is implemented in a first prototype, *ATL A Tutor Locator*. With it, we intended to test and validate our assumptions about ad-hoc transient communities and to validate the steps and selection criteria. We distinguish a request component, a population component, a community component, an indexing component, and a database to store relevant data (Kester et al. 2007) as depicted in Fig. 5.2. The request component covers the first two steps of the model (formulate and define context of support request) and provides the basis for the third step, identification of suitable peer tutors. The population component handles step 3 and the logistics required for step 4. The community component finalises the logistics and provides the facilities for the peers to enter the peer support process, up to the finalisation of the process (step 5). The indexing component precedes these steps and indexes all documents in the Learning Network. Finally the database stores personal data, data on learning activities and resources in the Learning Network, as well as results of the several components.

The Learning Network was implemented in Moodle, an open source platform for online learning environments (<http://moodle.org>). The four components have been developed as add-ons to Moodle or as an application that could be called from Moodle. The sources are made available under the BSD licence and can be downloaded from Sourceforge (<http://sf.net/projects/asa-atl>) or from our repository (<http://hdl.handle.net/1820/960>). For certain parameters in each of the steps we had to make decisions on how to implement and apply the model. These are described below.

5.2.3.1 Indexing Component

This component makes the preparations that are required for our model to operate on. First, it retrieves all resources that are used in the learning activities. These resources are split into better manageable text-fragments and saved as text documents. An index is created to record which document belongs to which learning activity; every document can occur in a single activity only. The documents constitute the text corpus that is required for LSA. The corpus is pre-processed by removing common words (stopping). LSA is then applied to the corpus to set up the term versus document matrix. This matrix is required to be able to match the question to the text corpus and find the documents to which the question correlates.

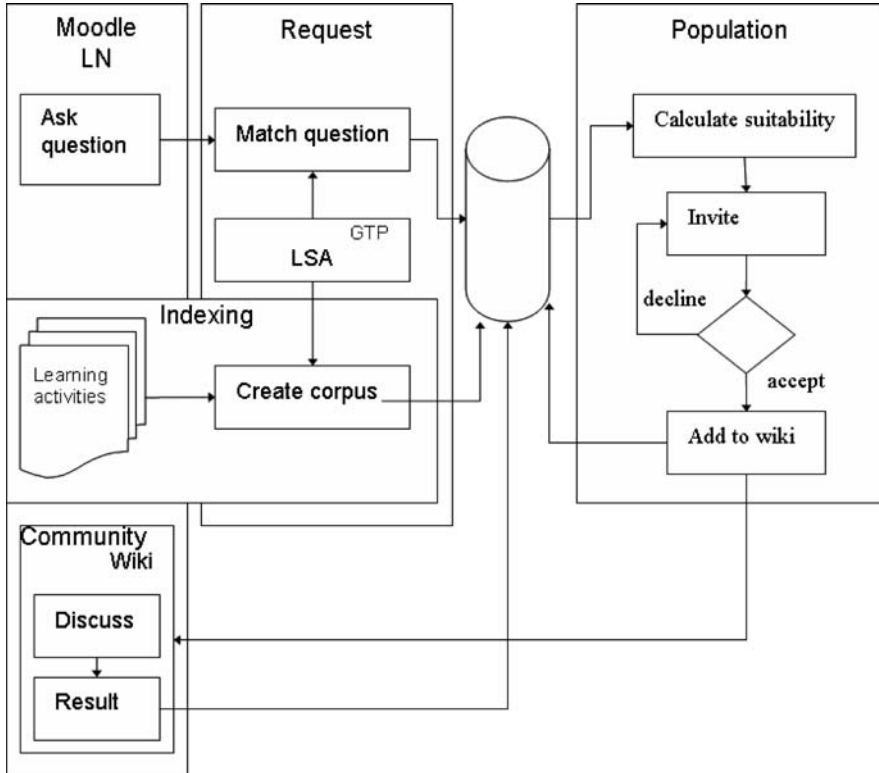


Fig. 5.2 The main components: Moodle constituting the Learning Network with learning activities, resources, question form and wiki; indexing component to create the corpus; the request component to define the context; the population component to select suitable peers and set up the wiki; and the community component for peers to discuss and formulate an answer (adapted from Van Rosmalen et al. 2008b)

We use the General Text Parser software (GTP) (Giles et al. 2003) to carry out LSA algorithms.

5.2.3.2 Request Component

This component provides the question form, assisting in formulation of the question, and determines the context of the question. The question form was developed as an add-on to Moodle. The members should be able to access the question form from anywhere in the Learning Network and from all learning activities in the Learning Network. The question form contains some instructions and guidelines for proper formulation of the question. There is some form of validation of the question, by checking whether the question is long enough. When the question is very short, a message appears, explaining that the question is not formulated well enough for the peer to understand what the question is about and the learner is prompted to refor-

mulate the question. The learner also has to indicate from which learning activity the question originates (if known) and has to provide details such as within which time limit an answer is required. Of course the learner identity is also required. In our case, identity is known automatically because the Learning Network requires authentication.

When the form is submitted several algorithms are invoked. First we verify the context of the question by using LSA to match the question to the documents of the learning activities. The documents with the top three (this number is configurable) highest correlations to the question are retrieved. Via the index in the database (as prepared in the indexing component), we can determine to which learning activity these documents belong. We then calculate a weighted sum of the correlation for each of these learning activities.

5.2.3.3 Population Component

This component is developed to select the most suitable peers to populate the ad-hoc transient community. There are two main processes, the selection process and the invitation process.

Selection process. In the selection process the content competence, tutor competence, eligibility and availability are calculated to return a rank-ordered list of suitable peers. The suitability is calculated as a weighted sum (value between 0 and 1) of the four criteria. The content competence is considered to be the most important and is included with a weight of 1; the eligibility and availability each with a weight of 0.5. In our situation, tutoring competence was not available, and thus received weight = 0. We calculate the other measures only for those learners with a positive value (> 0) for content competence. If the content competence algorithm fails, because there are no other learners with a positive competence for the relevant activities, we fall back to a ‘random’ selection in which we only take availability into account. The tutor suitability measure is a dynamic measure, and will be different every time the criteria are calculated. That is why the tutor suitability for every learner for every particular question is stored in the database.

Content competence. Preferably, content competence is based on assessment of competence and proficiency level attained. In our situation we did not have access to learners’ profiles nor could we conduct some form of competence assessment, because no data on competences were available. Most online learning environments however, offer the possibility to log members’ actions, including whether the learner has accessed certain learning activities. Moodle does not provide a profile, but logs whether learners accesses a particular learning activity. Every learning activity in Moodle can conclude with an assessment. Individual assessment grades can be retrieved from the database. In the present implementation the content competence algorithm is based on the completion status of the learning activity. As explained before, via LSA we can determine which learning activities are relevant for the question (via the correlation between the question and documents from the learning activities). This is required to determine content competence. For every learner we calculate the content competence for each of these learning activities, as a weighted

sum of the correlation and competence proficiency level for each of the learning activities. The correlation is provided by the LSA component. The content competence for each learning activity is determined by the score of an assessment, which is added to every learning activity. When the learner never accesses the learning activity, the content competence for that learning activity is set to 0; when the learner has started with the learning activity but not yet obtained a sufficient score for the assessment, content competence is set to 0.3; when the learner successfully passes the assessment, content competence is set to 1. A more accurate measure of content competence should take into account the score (or competence proficiency level), time expired since completion of the learning activity, and time required to study (time lapsed between start and successful completion).

Tutor competence. Tutor competence is a weighted sum of previous contributions, derived from logging and ratings by other peers. This is based on how actively the member participated in previous questions and how answers to previous questions have been rated.

Eligibility. Eligibility is a measure that indicates which peer is most near or similar to the learner asking the question. It is calculated over all learning activities in the Learning Network, but excluding those that correlate with the question. So the learning activity from which the question originates is omitted from the equation. For all other learning activities we need the content competence (0, 0.3 or 1) for every member. When the peer has the same content competence as the tutee, the value becomes 1, else 0. The total eligibility for the peer is the sum of all values divided by the number of learning activities involved. This is calculated for every peer.

Availability. Availability depends on presence and workload (e.g. currently to busy studying, preparing for exam, work), but should also take past workload (being involved in answering questions) into account. Data on presence can be taken from a calendar, but requires a calendar that members maintain accurately. Lacking this, we calculate availability as a relative measure on the basis of the number of invitations sent for the learning network, the number of invitations accepted by the member, the average number of invitations per member over a certain time period. This is set against a threshold M . A member can be at the left or right side of the threshold. Depending on where he is in the range, availability varies from 0, 0.25, 0.5, and 0.75 to 1. For a member who has contributed more than average, availability is set to 0 when the member contributed more than M above the average or to 0.25 when contribution is above average, but not more than M above average, while for a member who has not been very active, availability is considered to be high (1). The value for availability is set to 0.5 when the member contributed on average. We had to build in a certain period of time (set number of hours or days) over which to calculate availability to allow for members who did not respond (in time) and omit those for that set period from the selection algorithm. There are several reasons why a member does not respond to an invitation: not present, not available, time pressure, taking exams, holidays, etc. When insufficient members are available the invitation process fails. To prevent that these unresponsive learners are included in

new invitation cycles, we exclude those learners for a certain pre-defined period of time.

Invitation process. When the list of suitable peers has been calculated, the invitation cycle can start. First it checks whether the suitable peers have any outstanding invitations. To prevent a lot of invitations not being accepted because they get stuck in this loop, peers with outstanding invitations are removed from the invitation list. Invitations are sent out by e-mail to the peers selected, and contain the question as formulated by the tutee, the invitation to assist in finding the answer and instructions about how to proceed and accept.

A new page was added to Moodle to allow the peers to either accept or decline the invitation. When the peers accepted they are asked to provide some information on their alleged competence level.

The invitation cycle is complete when the specified minimum number of peers has accepted. In this case, another e-mail is sent to the tutee and peers to inform them that the invitation has been successful. The message contains the name of the peer who has accepted as well as instructions on where and how to access the ad-hoc community. It is possible that an insufficient number of peers accept, either because they decline or because they do not respond at all. After all, they might not be present when the invitation is sent. When none of the peers accepts, a new invitation cycle can be started, this time omitting the peers who declined from the selection list. Again, the list of suitable peers is checked against learners with outstanding invitations. When this cycle also fails, the tutee receives an e-mail message that unfortunately no peers can be found to assist with the question. The required group size (2), the minimum number (1) of peers that have to accept, the time period peers can take to respond (2 days) and how many invitation cycles (2) should be run, are configurable parameters.

5.2.3.4 Community Component

Finally, the ad-hoc transient community is set up. When the required number of people has accepted the invitation, a wiki is created. At that time, another e-mail message is sent to inform the tutee and peers that sufficient peers (in our case at least 1) have accepted the invitation and a hyperlink to the wiki is provided (A forum might be used instead, but is more suitable for discussion than jointly writing an answer). The wiki module of Moodle had to be modified to add instructions on how to use the wiki, guidelines on how to respond to the question, and to add the text of the question, and the three relevant documents. Access to the wiki is granted only to the tutee and selected peers.

A few more modifications to Moodle had to be made. In Moodle, wikis normally are added to a learning activity and are available to all learners. We had to provide an entrance to the wiki that is restricted to the members involved in a particular ad-hoc transient community. The link to the wiki is included in the invitation e-mail, but members need an alternative in case the e-mail gets lost. The tutee has to be able to end the process, and rate the contributions. In addition, all members require an overview of the status of the ad-hoc transient communities they are involved in.

Table 5.1 The peer selection algorithms

Explanation	Algorithm	Weights
<p>T_{SL} Tutor suitability of learner L</p> <p>Value between 0 (not suitable at all) and 1 (very suitable)</p> <p>Weights to adjust relative importance of the four measures</p> <p>Calculate only when content competence > 0, to assure a minimum knowledge level</p> <p>To assure presence, learners with available time = 0 are removed from the list</p>	$T_{SL} = \frac{(WT \times T_L) + (WE \times E_L) + (WA \times A_L) + (WC \times C_L)}{(WT + WE + WA + WC)}$	<p>WT = 0</p> <p>WE = 0.5</p> <p>WA = 0.5</p> <p>WC = 1</p>
<p>T_L Tutor competence</p> <p>Value between 0 and 1</p> <p>T_{W1} and T_{W2} to adjust the relative importance of T_e (on average how active the learner behaved in previous questions) and T_r (on average how previous answers were rated)</p>	$T_L = ((T_{W1} \times T_e) + (T_{W2} \times T_r)) / (T_{W1} + T_{W2})$	<p>Not available, since WT = 0</p>
<p>E_L Eligibility</p> <p>Value between 0 and 1</p> <p>E_L is relative to L_q (tutee)</p> <p>Calculated over all learning activities (AN) that do not relate to the question</p>	$E_L = \left(\frac{\sum_{i=1, \dots, N \text{ \& all } j \text{ANI is not question related}} (\text{score}(\text{ANI}_{L_i}))}{(N - \# \text{question} - \text{related ANI's})} \right) /$	<p>AN score can be:</p> <p>0 – not started</p> <p>0.3 – started</p> <p>1.0 – assessment completed successfully</p>

Table 5.1 (continued)

Explanation	Algorithm	Weights
<p>A_L Availability Value between 0 and 1 M (max_extra_workload) and T_p (time period over which workload is calculated) Depends on relative past workload, comparing number of times a member is involved relative to other members in given time period</p>	<p>$A_L = \text{one of } [0, 0.25, 0.5, 0.75, 1]$ 0.5 if L contributed on average; 0.25 if L contributed above average but not more than M above average; 0 if L contributed more than M above average</p>	<p>$M = 1$</p>
<p>C_L Content competence Value between 0 and 1 D: number of documents to calculate LSA correlations for D_i: number of documents provided W_{ANI} based on correlation between question and documents in ANi C_{ANI}: content competence for ANi, taking into account score, time expired since completion and study time of ANi</p>	<p>$C_L = \frac{(W_{ANI} \times C_{ANI}) + (W_{AN2} \times C_{AN2}) + \dots + (W_{ANn} \times C_{ANn})}{(W_{ANI} + W_{AN2} + \dots + W_{ANn})}$</p>	<p>$D = 3$ $D_i = 3$ C_{ANI} based on score of ANi</p>

Adapted from Van Rosmalen et al. (2008b).

This was accomplished by adding a couple of additional pages to Moodle. The first page provides the member with an overview of the status of the ad-hoc transient communities they are involved in, either because they invoked them or because they accepted an invitation of somebody else. Via this page, the members can enter the wiki. The tutee also can change the status of the wiki to 'Completed'. When the tutee changes the status, he is asked to rate the contribution. After the status has changed to 'Completed' access to the wiki is closed, but the wiki and all contributions are stored.

5.3 Conclusion

Applying the peer tutoring ad-hoc transient communities in an actual Learning Network provided promising results. The model was evaluated in an experimental setting (Van Rosmalen et al. 2008b), involving over 100 learners divided in two groups. In the experimental group, the following weights applied: 1.0 for content competence, 0.5 for availability, and 0.5 for eligibility. In a newly started Learning Network, learners have not yet developed any content competence. Because content competence is an imperative constituent of the model, this could result in 'cold start' problems: the selection does not return any suitable persons because none has the required content competence. To prevent that from happening, we resort to the same selection mechanism as used for the control group. In the control group, we only tried to distribute the questions evenly across all members by taking availability and past workload into account. Thus, the selection criteria were applied with weight 1.0 for availability and 0 for all other criteria.

We looked at the number of questions solved successfully. This was determined by the learner asking the question and by two expert tutors. For the system to be effective at least 50% of the questions had to be solved by the peer tutors. To determine whether our selection criteria made any difference, we compared the number of invitations accepted, the time needed to answer the question, the number of questions answered and the quality of the answers in the experimental group with those in the control group. We also looked at whether the peers used the three documents and whether they perceived them to be useful.

More questions were asked and above all a much larger percentage of the questions was answered in the experimental group (71%) than in the control group (45%). More importantly, a single invitation cycle was sufficient in around 80% of the questions in the experimental group, while in the control group at least 2 invitation cycles were required for at least half of the questions. In addition, in the experimental group questions were answered within a day and a half, while the peers in the control group took more than 2 days. In the control group, a larger proportion of the learners was involved, as side-effect of the model that tries to distribute workload evenly among the learners when no other measures are available.

Initially we allowed peers a period of 6 days to respond to an invitation. Due to the limited number of members in the Learning Network and the small number of invitations being sent out, this resulted in many outstanding invitations not being

answered, while the same members were repeatedly invited for multiple questions. Therefore we had to decrease this response time to 2 days to prevent the selection from failing. We did not have proper availability data, other than past workload. Having members indicate the time and days when they are available, or determining when somebody is online can alleviate this problem.

The peers were asked to rate their own competence in answering the question when they accepted or declined an invitation. And the tutee rated the answer. In the experimental group, the peers were more outspoken in their ratings: more either absolutely accepted or declined, while in the control group the ratings were more evenly distributed. This indicates that the selection criteria returned the more suitable peers.

Although the supplied documents were used in a small proportion of the questions, members tended to rely more on course content, prior knowledge or non-specified resources; the learners in the control group used the documents more often. In particular, when asked to rate the usefulness of the documents they rated them at a higher level. Apparently the documents were more useful to the learners in the control group as the learners were randomly selected. The usefulness of the documents also depends on the type of question, the complexity of the domain and the previous knowledge and content competence of the peers. When the selection delivers a good match between question and suitable peers, the documents might not be needed.

There were some drawbacks to the implementation. The participants were not familiar with wikis and did not engage in a joint writing process. Instead they used the wiki more like a forum. Care should be taken to either use a facility participants are familiar with or to provide elaborate instructions on its purpose and proper way of use. The wiki did not notify peers when the wiki was changed; so some members complained that they wanted to elaborate on their question but no longer got a response from the invited peers. Nor were the people involved notified when the tutee closed the question or when it was successfully answered. A clear notification system seems imperative. People at least need to be able to see what is changed and preferably are notified about any changes. The peers in the ad-hoc transient community need additional communication facilities, for example by providing e-mail or Skype contact details. These kinds of data can be added to a personal profile the learners can share. In the personal profile the learners can indicate what information is public and what is private.

As discussed, the model was tested in an experimental settings and the learners were aware of this. This might have prompted a more extensive use of the system. The experiment ran only for a short period of time, while a Learning Network won't have a fixed duration and learners will not all start and end at the same time. Additionally, the Learning Network domain was rather small and of a relatively simple level. In more real-life Learning Networks incentives for participation should be made clear and policies such as described in Chap. 3 should be applied to ensure prolonged participation.

Interestingly and importantly, all members indicated that answering questions is a good investment of time. The mostly selected reasons given were: 'I'm aware

that other students also have questions' and 'It improved my knowledge and understanding'. In particular the learners in the experimental group gave more positive feedback and indicated that they would like to use ad-hoc transient communities in other Learning Networks (Van Rosmalen et al. 2008a). Furthermore the members remarked that they wanted facilities that allow them to contact the peers outside the context of the peer tutoring ad-hoc transient community. This can be taken as a first step towards enhancing social interactions in the Learning Network as a whole.

We can conclude that ad-hoc transient communities are well-suited for answering content-related questions by involving peers. Moreover, the model implementation offers sufficient evidence that ad-hoc transient communities contribute to improving the community aspect of the Learning Network by enhancing the social interactions. The most important aspects to be covered by the service are:

- Base the peer selection criteria should at least be based on eligibility and past workload. Important for the eligibility measure is the content competence relative to the other members.
- As the eligibility measure depends on an accurate measure of content competence, record data for this in a learner's profile; logging at least measures of content competence, e.g. completion of activity, date activity started and completed.
- Ensure that the question can be related to the correct activity or competence by using appropriate technologies.
- Allow learners to decline only a set number of invitations.
- Obtain valid availability data, e.g. by using a calendar.
- Allow members to contact the peers outside the context of the ad-hoc transient community, e.g. by adding the peer to one's list of contacts, or friends, noting reasons for doing so.

References

- De Vries, F.J. et al.: Identification of critical time-consuming student support activities in e-learning. *ALT-J Res. Learn. Technol.* **13**(3), 219–229 (2005)
- Giles, J.T. et al.: GTP (general text parser) software for text mining. In: *Statistical data mining and knowledge discovery* (CRC Press, Boca Raton 2003) pp 455–471
- Kester, L. et al.: Matchmaking in learning networks: Bringing learners together for knowledge sharing. *Interact. Learn. Environ.* **15**(2), 117–126 (2007)
- Landauer, T.K. et al.: Introduction to latent semantic analysis. *Discourse Process.* **25**, 259–284 (1998)
- Van Bruggen, J. et al.: Latent semantic analysis as a tool for learner positioning in learning networks for lifelong learning. *Br. J. Educ. Tech.* **35**(6), 729–738 (2004)
- Van Rosmalen, P.: *Supporting the tutor in the design and support of adaptive e-learning* Unpublished PhD thesis (Open Universiteit Nederland, Heerlen 2008)
- Van Rosmalen, P. et al.: Knowledge matchmaking in learning networks: Alleviating the tutor load by mutually connecting learning network users. *Br. J. Educ. Tech.* **37**(6), 881–895 (2006)
- Van Rosmalen, P. et al.: A model for online learner support based on selecting appropriate peer tutors *J. Comput. Assist. Learn.* **24**(6), 483–493 (2008a)
- Van Rosmalen, P. et al.: A learner support model based on peer tutor selection. *J. Comput. Assist. Learn.* **24**(1), 74–86 (2008b)

We argue how this can be achieved by providing Personalized Recommender Systems (PRS) that provide personalised advice on the best next learning activity to study. This section on navigation services in Learning Networks aims to translate some ongoing research into considerations, practical guidelines and concrete examples for practitioners that are considering to set up similar navigation support in their Learning Networks. This will be achieved along four lines of study.

In Chap. 6 we argue which combinations of recommendation techniques will provide best support for navigation (or way-finding) in various situations of informal and formal learning. Using hybrid recommendation techniques for Personalised Recommendation Services (PRS) appears to be a novel and promising approach to set up individualised navigation in Learning Networks. Such an approach combines the collaborative filtering of information about others' behaviour with matching information connected to individual learners and activities.

In Chap. 7 we present guidelines on how to implement a PRS in a concrete Learning Network (i.e., in the domain of Introductory Psychology). The chapter discusses and presents techniques and concepts that have to be considered when setting up such an experimental study, offers an overview of suitable variables in an evaluation framework to evaluate the effects of the recommendations, explains how the experimental data have to be prepared and analysed, and presents an overview of available recommender systems.

Chapter 8 describes the potential and considerations of using simulation studies to define the requirements for PRS, without having to take into account various practical real-life constraints. A conceptual simulation model will present most relevant variables, relationships, and formulas of such PRS. The simulation results guide further articulation of the conceptual simulation model and ultimately provide guidelines for future PRS-system development. The simulation study reveals that a rating-based light-weight PRS is a good alternative for ontology-based recommendations, in particular for low-level goal achievement.

Finally, Chap. 9 explains how a learning path specification could be used to help learners in finding their way through a Learning Network. A learning path is a set of one or more learning activities leading to a learning goal. In order to support learners in comparing and selecting suitable learning paths in a Learning Network, a formal and uniform way to describe learning paths is needed. For that purpose a learning path specification has been developed. The last chapter of this section explains how the learning path specification facilitates these processes. It also explains how the specification can be used to describe both formal and informal learning paths.

All chapters in this section on navigation provide valuable information for both researchers, professional developers and technical designers. Chapters 6 and 9, because of their more conceptual nature appears most useful for researchers. Chapters 7 and 8 provide concrete information on actually setting up services and simulation studies and seem most useful for technical designers.

References

- Darken, R.P., Silbert, J.L.: A toolset for navigation in virtual environments. *Proceedings of the 6th annual ACM symposium on User interface software and technology* (pp. 157–165), (1993), Atlanta, Georgia, USA: ACM Press

providing (individualised) navigation services. In order to provide the reader with the necessary overview, we will present a model of navigation in Learning Networks. The model will make clear how PRS should be related to learning technology specifications required, like the learning path specification, the learner profile description, and the competence description, as well as to possible recommendation techniques, like information-based and social-based techniques. Every reader will immediately recognise problems with way-finding on the Internet in general. The accelerated growth of the World Wide Web has turned the Internet into an immense information space with much, very diverse and often poorly organised content. Successfully searching and finding information on the Web turns out to be hard, especially for novice users. Results with secondary school children for instance indicated that experienced WWW-users are more proficient in locating relevant information than novice WWW-users (Lazonder et al. 2000). The observed differences were ascribed to the experts' superior skills in operating web search engines. Hölscher and Strube (2000) found that online users need additional Internet skills training for query-based searching and intersite navigation. Other researchers have developed a variety of tools (like a task-specific portal to structure information, worksheets for reflective web use, process-worksheets, prompts, and discussion sessions) that help regulate the search process (Brand-Gruwel and Gerjets 2008).

In distributed Learning Networks (Koper and Sloep 2003) online users are confronted with similar problems in finding information, i.e., in navigating towards the most relevant learning activities. They are populated by groups of professionals that usually are more heterogeneous than in regular education: various ages, needs, motivations and preferences characterise the professionals. Individual professionals therefore will need individualised navigation support. The problem of finding information can haunt professionals during their stay in dynamic Learning Networks, especially when learning activities may be added or withdrawn and where learners may change their learning plans. This problem is aggravated if the learner also has to choose from the offerings of various content providers, that might vary strongly on several aspects. Navigation services will then be needed to help the learner find and select most adequate activities. Personalised Recommender Systems are meant to enable individualised learning paths that lead the learner through an effective study progress.

Because of the fact that information in Learning Networks is always aimed at learning towards a certain goal, such navigation services have to take into account additional pedagogical issues (e.g., activities have to build on the existing competence level of the learner, dependencies on prerequisite activities and practical constraints have to be taken into account). Compared to looking at interesting news or movies, or to listening to preferential music tracks on the Web, learning activities cannot be studied in any order, they have to be carefully sequenced.

Most curricula have been designed carefully by professionals in the field. It has been learners' main task to follow the sequence that was designed in the curriculum. It is highly questionable whether available curricula contain the most suitable order for each learner. Besides, from a professional development perspective, learners will not just follow available curricula like is the case in more regular education.

At various stages of their lives they will face the task to select and mix both formal and informal learning activities, taken from different providers and sources (like the Internet, peer discussions, or training courses). In the absence of any ‘designed order’ of a curriculum, it will be hard for learners to select and sequence the right learning activities. Learners’ problems in ‘way-finding’ (the process of selecting and sequencing learning activities, which we consider synonymous to ‘navigation’) will decrease the efficiency of education provision (the ratio of output to input) and increase the costs. Moreover, traditional institutional facilities like course catalogues or face-to-face study advice do not offer adequate guidance within the context of distributed networks for professional development, in which learners have to make well-informed choices from the vast amount of learning activities offered by different providers such as institutions, tutors or peers.

Although research reveals a relation between advice and drop-out rates, advice appears to be just one of many factors (Rovai 2003). There may be other alternatives for costly face-to-face advice (e.g., domain specific diagnostic self-tests), but this section will focus on Personalised Recommender Systems (PRS) Personalised Recommender Systems (PRS) that advise learners on suitable learning activities and paths towards certain learning goals. We focus on a specific type of PRS based on a combination of collaborative filtering information about the behaviour of other learners (social-based approach) with information about learning activities and needs and preferences of learners (information-based approach). These hybrid systems provide learners with individualised navigation services that advise suitable learning activities and paths towards certain learning goals, like the attainment of competences. The combination of the social-based approach with the information-based approach has hardly been applied in learning (Herlocker et al. 2004).

This first chapter of the navigation section will provide you an overview of the main navigation concepts involved and the main questions that need to be resolved when applying this combined approach to Learning Networks. Important practical issues that will be addressed are: what are the most adequate recommendation techniques and how can they best be combined into recommendation strategies, what are the most relevant elements and attributes (from learners and learning activities) to be taken into account in the PRS, and how can we assure interoperability between learning activities and learning paths.

The second section of this chapter provides a short introduction to relevant work, relevant learning technology specifications (uniform descriptions of competence, learner profiles, and learning paths), and existing recommendation techniques. The third section presents a preliminary model of the elements and attributes involved in navigation services. An important guideline here is that PRS should not rely on rather specific and intensive data provisioning and maintenance. Finally, the fourth section arrives at some conclusions and suggestions for future research into navigation services.

The character of this first chapter has to be of a more conceptual nature in order to better understand the remainder of the section. The character of the next three chapters will be more practical. The next chapter, Chap. 7, of this navigation section provides some concrete guidelines on practical issues and choices when actually

setting up a concrete PRS in a Learning Network, with examples from an experiment in the domain of Psychology. Chapter 8 of this section describes practical guidelines on how simulation studies can help you in designing more light-weighted PRS that will enable more effective, more satisfied, and faster goal attainment. Chapter 9 of this section will then propose a concrete learning path specification that could facilitate the navigation process.

6.2 Learning Technologies for Personalised Recommendations

We start this section by introducing related work on personalised recommendation (Sect. 6.2.1) from other domains. Then some guidelines on using relevant learning technologies (Sect. 6.2.2) and recommendation techniques (Sect. 6.2.3) to enable such personalised recommendation in Learning Networks are provided.

6.2.1 Relevant Work

Most readers will be familiar with recommender systems that offer advice when looking for books, movies or music (e.g., amazon.com). Such applications are based on collaborative filtering of information obtained from the behaviour of other buyers on the web (others that bought this book, also bought these books). Some even take ratings or tagged interests by individual users into account (others that like Tarantino as director, also like these directors). Review studies do not mention such personalised recommender systems in learning (e.g., Brusilovsky 2001; Burke 2007; Herlocker et al. 2004). This probably is the case because the educational field imposes additional and specific demands on the advice required. Main differences between selecting books for reading and selecting learning activities for study are the degree of voluntariness and required order (as most learning activities are required to obtain some learning goal or prerequisite to another), and the possibility to establish an explicit completion (as most learning activities are to be assessed for successful completion). Such differences impact learner's motivation, and the way personalised recommendations for learning activities should be provided. Preliminary explorations in the e-learning domain are yielding promising results, and by setting up simulation studies you can gain more insight into the additional dynamics of Learning Networks.

A simulation study of a recommender system to support learners in a Learning Network has been reported by Koper (2005). He simulated rules for increasing motivation and some other disturbance factors in Learning Networks, using the Netlogo tool. Learners had to complete a certain set of learning activities, and after each completion were 'set' to complete the best next learning activity, based on the successful completion of next learning activities by others. Amongst other factors, the provision of this indirect social navigation accounted for about 5–12% of the increase

in goal attainment (completion of the set), depending on the ‘matching error’. An interaction effect showed that recommendations can compensate for bad matching.

Related to this simulation study are some real-life experiments we carried out, as have been reported by Janssen et al. (2007) and Drachsler et al. (2008). The aim of the Janssen study was to recommend most efficient learning paths (like with the ‘shortest route’ provided by the GPS in your car) in general, leaving the learner and learning activity characteristics out of scope. The Drachsler study (for further information, see Chap. 7) also explored which paths are most *attractive* or *suitable* for individual users (like with the routes suited for bicycling). Authors offered learners a similar recommendation (Most successful learners continued with Y after having completed X). The recommendations in the first study did not take personal characteristics of learners (nor possible ‘matching error’) into account. In the second experiment collaborative filtering information was combined with personal information taken from a user profile, an hybrid approach that appeared to solve some common problems when using collaborative filtering only (e.g., the cold start problem). In both studies, the indirect social navigation appeared to enhance effectiveness in a Learning Network (completion of the set of learning activities), but only in the second study did it increase efficiency (the time it took to complete them).

Although these first results of applying recommender systems for sequencing learning activities in real-life Learning Networks appear promising, significant effects appeared to be still relatively small and to depend on the availability of sufficient amount of learners and elapsed time. A promising way to further explore possible improvements is by means of simulation studies. Where ‘matching error’ was limited to competence levels of learners and learning activities in the Koper study, mismatching can also occur on other learner characteristics such as personal needs, preferences and circumstances. In another simulation study by Nadolski et al. (2008), we did not only take the competence gap but also the preference gap and some other learner and activity characteristics into account (for further information, see Chap. 8). In this respect, Bocchi et al. (2004) found that learner characteristics accounted for about 30% of retention rates in an online MBA program. Personalised recommender systems filter specific data from learning activities and learners that fit individual needs, interest, preferences, constraints or circumstances. Results show such an hybrid approach to be most effective. Relevant recommendation techniques to be combined in such an approach will be described in Sect. 6.2.3 of this chapter. We will now first turn to relevant learning technology to be used for navigation services.

6.2.2 Learning Technology Specifications

Learning technology specifications are needed to enable interoperability, reuse and exchange of learning objects and actions. We discuss three of such specifications that are closely related to navigation services and personalised recommendation in a Learning Network for a certain domain: learner’s start position in that domain (based

on learner's prior learning history or *e-portfolio*), the aimed *competence* profile for that domain (also called the competence description), and the *learning path* from already acquired competence towards the acquisition of new competence. Way-finding or navigation services support selecting and sequencing available learning activities into an individualised learning path. A learning path description is currently non-existent but needed to uniformly describe different possible combinations of learning actions that lead to certain learning outcomes. E-portfolio and competence specifications do exist but need elaboration to sufficiently cater for personalised recommendations.

Therefore, personalised recommendation in a Learning Network needs to combine and use following pieces of the learning technology puzzle (some to be further addressed in the remainder of this and consecutive chapters): Uniform and meaningful description of (both formal and informal) *learning paths*; *Learning activities* that are addressable and meaningfully described; Uniform *e-portfolios* that define needs and preferences; Uniform *competence descriptions* that define proficiency levels; A learning path processing engine able to compute what remains to be done by the learner to acquire the competence profile; An engine recording completion of activities and propagates this to associated systems; and Information *matching techniques* to enable personalised recommendations. To exchange competences and learning activities and -paths in an interoperable way, we will first need learning technology specifications describing these concepts in a both uniform and meaningful way. Some descriptions and future directions are provided below.

Learning path description. If all providers would use a common language to describe their learning programs and activities, personalised recommender systems could better support learners deciding between various paths to reach aimed competences. When the learning paths taken by predecessors could be stored in a commonly used format, novice learners could benefit from the choices made and experiences gained by more expert learners that have already acquired the competence, by finding and following the paths of (similar) peers. But unfortunately, no such commonly used format exists. A number of existing approaches to (formal) curriculum modelling (e.g., CDM 2004; XCRI 2006) do work in this direction. Tattersall et al. (2007) propose IMS-LD (2003), a specification for modelling learning designs, as a strong candidate to model learning paths as well, and demonstrate that its selection and sequencing constructs appear suitable on both the level of learning activities (units-of-learning) as well as on higher levels of granularity (like professional competence development programmes). They note that, in addition to the curriculum structuring concepts covered by IMS-LD, other information will be required to provide learners with more personalised advice on learning content.

Janssen et al. (2008) have built on this approach and propose a number of key elements to describe learning paths relevant to navigation on a basic level. Besides learning activity, learning objective and prerequisite, these elements include structure (for grouping dependent learning activities), method (for ordering learning activities into a scenario), conditions (to model the order to accommodate personalisation), and metadata (including additional criteria). Additional criteria could include, for instance, information about study costs, required study time, mode of

delivery, teaching place, guidance, way of assessment, etc. This information will be helpful in achieving more personalised overviews of available learning paths. Chapter 9 of this section will provide some concrete guidelines when using this learning path specification for navigation purposes.

E-portfolio description. Describing and recording learner's history and profile becomes of crucial importance for professional development. Competences are not only attained through formal education, but also through work, at home or in any other context where problem solving takes place. Currently, educational providers have little possibilities to adapt their formal programmes, and to take into account individual needs and preferences or prior knowledge attained through other (formal) education. It becomes even harder for them if they are asked to take into consideration competences acquired through less formal or informal learning. Uniformity in assessment (to measure and accredit) and e-Portfolios (to store) to enable this in the future constitute a research field on its own. The IMS Learner Information Package (IMS-LIP 2001) and ePortfolio (IMS-ePortfolio 2004) specifications ensure exchange of learner records, by linking to produced artefacts and formal achievement records like references. However, fields describing learner information are open and optional, so these specifications do not provide classifications of specific learner information that is easy to interpret uniformly.

Berlanga et al. (2008) have drawn up some additional requirements for an e-portfolio description to be useful for more personalised navigation in Learning Networks. They have argued that e-Portfolios should allow learners to present themselves, to reflect on how they acquire competences, and to show and manage their social presence in the Learning Network. To this end, e-Portfolios should integrate, respectively, rhetorical, pedagogical, social, and technical perspectives. The rhetorical perspective is needed to show the learner's competences mastered, including a history of the professional competence development programs, and learning activities the learner has followed, an overview of communities the learner has been involved in, and the showcases she has created. The pedagogical perspective aims at supporting learner's self-reflection, so the learner can define the competences she masters, review and create (new) professional competence development programs, and assess her competences. The social perspective aims at fostering interaction (by creating and maintaining a personal profile as well as contacts within the Learning Network) and social help support (by defining mechanisms for peer-support). The technical perspective is needed to support the other three perspectives.

Competence description. Cheetham and Chivers (2005) define competences as the integrated application of knowledge, skills (or competencies), experience, contacts, external resources and tools to solve problems at a certain level of performance in a certain occupation or any other context. According to this definition, competence is related to three dimensions: the *type* of competence (i.e., cognitive, functional, personal, professional or ethical); an occupation or performance *context*; and the proficiency *level* of a person with respect to an occupation or context. The IMS Reusable Definition of Competency or Educational Objectives (IMS-RDCEO 2002) and the IEEE Reusable Competency Definitions (IEEE-RCD 2006) specifications do not have much semantic value, since they simply attach IDs for a registry

and URIs, and only reference to more controlled models. The HR-XML consortium (2006) goes one step further in the description of a competence and includes classes (as free text fields) where ‘evidence’ (with reference to learner portfolio) and ‘proficiency level’ (complexity, intensity, quantity) *can* be described.

We suggest that a competence description should at least include an interoperable classification of *type* of competence and proficiency *level*, and preferably also have a classification of *context*. Sicilia (2005) proposes the formalism of ontologies to express more details in competence schemas, in order to be connected to learner profiles and learning activities. Matching learning activities to the right learners requires ontological structures and more meaningful descriptions of competences in a registry (Ng et al. 2005). Prins et al. (2008) have stressed the importance of identifying the complexity of learning and performance situations for personalised recommendation, and propose to extend current competence definitions with a context-concept, named a ‘learning and performance situation’ (LP-situation). This context-concept will further be determined by values on three factors: a complexity factor, a proficiency level and a performance indicator. Inclusion of such a context-concept will also help defining the proficiency level and to design learning tasks and CDPs. Authors propose to use various (local) schemas for defining these determining factors, but to have one mapping function towards uniform proficiency levels. This way we could find a compromise between maximal flexibility to individual providers to define their own factors and values, and a uniform mapping mechanism to enable a more meaningful exchange of competence definitions.

Closely related to competence definitions is the concept of a *competence map* that defines the relations between various competences. It could be questioned whether there should be such maps for more informal Learning Networks since the professional should decide on the most suitable order and preferred structure. Assuming that such a predefined structure is desirable and possible, several research questions need to be answered that pertain to the hierarchical organisation in the map, to aligning particular customs for structuring certain domains, and to connecting such maps to navigation (and other learner support) services.

6.2.3 Recommendation Techniques

Let’s for now assume that all required information about learners, targeted competences, and required learning activities is available in an exchangeable format. We will then need to decide upon techniques to match information about learners with most suitable activities. There are a number of drawbacks and advantages to current techniques that will be addressed in this section (Chaps. 6, 7 and 8). Our approach is to look for a hybrid approach, combining the advantages of both information- and social-based information techniques. After explaining the concept of information matching, we will describe those two types of techniques. A practical, stepwise tutorial on how to implement an hybrid recommender into a concrete, self-organised Learning Network can be found in Chap. 7 of this book section.

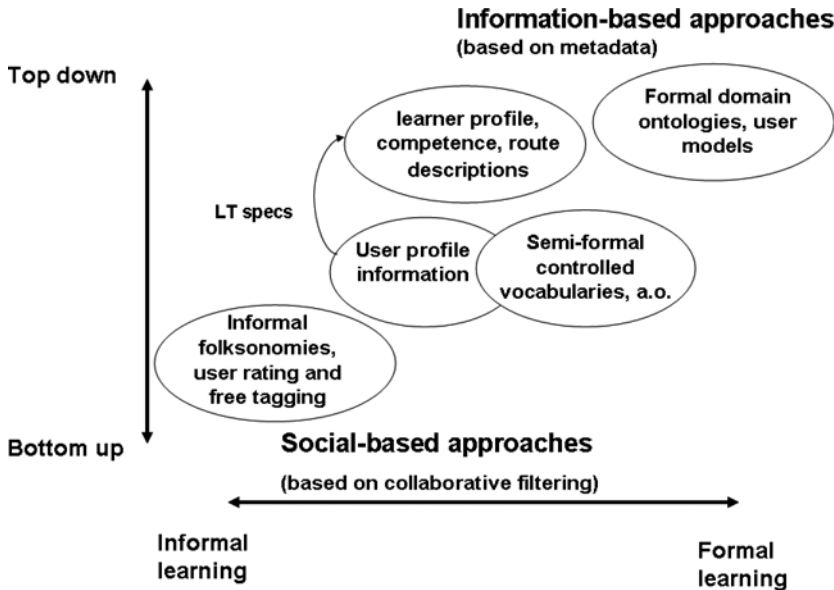


Fig. 6.1 Social-based versus information-based approaches

Information matching or recommendation techniques work with two entities: users and items. Elements of both entities are being associated with a profile carrying certain characteristics. The utility of an item (i.e., learning activity) to a user (i.e., learner) is usually represented by a rating function $R: Users \times Items \rightarrow Ratings$. Recommendations are estimated ratings for items which have not been ‘seen’ (i.e., enrolled, rated, successfully completed) by the user. Within information matching techniques a distinction is made between *information-based approaches* (based on learning technology standardisation, metadata and semantic web efforts) and *social-based approaches* (based on data mining, social software, and collaborative filtering) (Balabanovic and Shoham 1997; Van Setten 2005).

Figure 6.1 depicts the relation between the two approaches, the type of learning (formal versus informal) and the formalisation of information in terms of learning technology (imposed from the top downwards versus emerging from the bottom upwards).

Information-based approaches. Information-based approaches may use certain keywords or metadata which represent certain characteristics of the learners and learning activities. The system then keeps track of items the user was previously interested in, and recommends items with similar or related keywords or metadata. Similarity of items is calculated with techniques based on item-to-item correlation that may use keywords in documents (Schafer et al. 1999).

Ontological modelling has the immense (potential) advantage of exact matching competence descriptions in the learner profile (user model) with available professional competence development programmes (domain model), sharing common

understanding in a machine readable way. Exact matching is often required in more formal learning with, for instance, more discrete criteria to allow students to enter or progress than in more informal learning where more fuzzy matching might suffice.

A serious *drawback of modelling* is the enormous amount of work in enriching resources with metadata and the arbitrary character of such models. For example, main problems in instructional planning are caused by limitations to *learner* modelling, for instance because prior knowledge can not be uniformly compared for all learners. A second limitation is the inability to analyze and capture all important characteristics of the *content* (e.g., How to distinguish an excellent from a poor article that is using the same keywords?). Most categories we employ in life are based on fuzzy concepts rather than on objective rules, e.g., Where's the line between good and excellent? Do we really care about the subtle distinctions in wine-tasting? (Morville 2005). A third drawback is caused by the fact that words in the form of metadata are ambiguous. Our language is filled with synonyms, homonyms, acronyms, and even contronyms (words with contradictory meanings in various contexts). In retrieval the forces of discrimination and description are battling, with full-text search being biased towards description (finding general words with many meanings) (like with Google; see Brin and Page 1998), and unique identifiers (like an ID number of each competence in a registry; see IMS-RDCEO 2002) offering perfect discrimination but no descriptive value whatsoever.

It will neither be possible nor needed to fully model Learning Networks, but certain levels in formalising information can be distinguished. *Ontologies* describe how concepts of the world are related and represented using formal relations. An ontology is a rather strict formalisation into a machine-readable format consisting of entities, attributes, relationships and axioms (Guarino and Giarretta 1995). There are also semi-formalisations holding the middle ground: *Structured meta-data* fields (with headers like title, domain, provider of the activity), *faceted classifications*, which permit to assign multiple classifications to an object and to accommodate way-finding that varies by user and task, and *controlled vocabularies* (e.g., fixed categories, keyword lists, audiences) that try to control the language ambiguity.

Social-based approaches. The big advantage of social-based approaches is that they are completely independent of the representation of knowledge in domain or user models. Instead of recommending to a specific user the items similar to previously liked items, this approach recommends the items liked by other users in similar situations or with similar preferences (peer groups). It uses peer opinions to predict the interest of others, and matches users against the database to discover historically similar tastes. It avoids the enormous amount of work in enriching resources. For instance, Li et al. (2005) combine item-based and user-based collaborative filtering, based on the content information and ratings at the same time, which make it possible to alleviate both sparsity and cold start problems. For instance, when a database of learning activities would contain both the ratings of

peers (user-based) and tags describing the pedagogical taste (item-based, e.g. programmed instruction or problem-based learning), it would become possible to personally recommend items targeting the taste of users which will be apparent from their history of rating.

A serious *drawback of recommendations based on social-based approaches* is their limited value for new or few users, that is when it is hard to find similar users or when just a few users have rated the same items, or when no content is available about already attained competences. Clustering then may be an alternative to solve this sparsity problem (Agarwal et al. 2006). Social-based recommendation techniques also suffer from serious limitations when exact matching is required, and when competences in domains go beyond the verbal realm (e.g., hard to express communication or motor skills). Although, for instance, collaborative filtering has been considered as a mainstream social-based technique for recommender systems, applications with actual learning behaviour data for recommending learners have remained scarce.

There also are more technical limitations to collaborative filtering. Collaborative filtering relies on large numbers of users explicitly rating or completing learning activities. When recommendations would solely depend on collaborative filtering, new or few learners would be seriously handicapped by ‘cold-start’ problems. Furthermore, the events registered in common log files format (as defined by W3C) are extremely low level, which complicates further analysis and more clear cut decisions that might be required in more formal learning curricula (e.g., only learning activities of certain types are allowed, or only grades above a certain threshold are considered sufficient for their completion). This makes it difficult to know which (type of) users are interacting (since only IP addresses are logged) and what (type of) interactions they are engaged in (since only URLs are logged).

One possible solution to (partially) address these limitations is to enhance the logs with additional information drawn from semi-formal descriptions about which learner did what and whether this was successful (Oberle et al. 2003). Especially with small amounts of learners, more information is required for exact matching of learners and activities in formal learning programs, and more information-based (or even ontology-based) approaches will come into play. Using more structured, but less formalised methods to describe learner profiles, competences and learning activities (like structured metadata or controlled vocabularies), might offer a more feasible and intermediate position. Such structured descriptions can be the basis for standardisation through learning technology specification (see learner profile, competence and route descriptions, Fig. 6.1). Another improvement could be to not only analyze the collective, *average* behaviour of peers, but also take individual, *personal experiences* as a starting point. This way not only suitable learning activities, but also suitable peers to study with could be advised. In summary, using Learning Technology specifications, and combining social-based with information-based approaches to match learners and learning activities, might be an approach to way-finding that addresses the cradle-to-the-grave challenge posed by both formal and informal professional development.

6.3 Modelling and Experimenting with Navigation Services

As described before, personalised recommendations should support the learner to *compose most suitable sequences of learning activities* to attain competences. A model for navigation services has been proposed (Sect. 6.3.1) and a study using a personalised recommender that was implemented according to this model will be described (globally in Sect. 6.3.2, and in more detail in the next Chap. 7).

6.3.1 Model for Navigation Services

Our model for navigation services in Learning Networks is depicted in Fig. 6.2. The model shows the elements (classes) and values (attributes) that play a role for navigation services. The model is still under development since concrete attributes of the classes are to be decided and validated (denoted by the ‘...’), as well as the interfaces with other services and learning technology specifications. The model intends to describe relevant classes for providing personalised recommendations. The model builds on and partially extends the domain model developed in the TENCompetence project (for direct reference, see <http://hdl.handle.net/1820/649>) which is also described in Chap. 18 of this book. This section does not explain the (rather complex) model in detail but just describes its main classes relevant here, with illustrations from a use case example which will be elaborated in the next chapter. Our use case deals with a professional interested in finding a personalised learning path in a Learning Network of psychology courses on an introductory university level. According to our model for navigation services (classes between square brackets), an [Actor] (which role can be learner) has a learning [Goal] (of the type ‘competence’) that make her perform [Actions]. Characteristics of an actor (like study time, study motive, study interest) can be modelled according to the *IMS-ePortfolio* (2004) specification. [Actions] can be aimed at the acquisition of certain [Competences].

In our use case, novice psychology students were provided a Learning Network with a collection of 17, largely independent courses on various topics on an introductory level. Learners were allowed to find and select the most efficient learning path through this collection, depending on their individual needs and preferences. When action are completed with an [Assessment Result] available, one can infer whether the [Competence] is mastered at a certain [Proficiency Level].

Actions can include learning paths (named [CDP]), [Units of Learning] (modelled in IMS-LD), [Activities] (not modelled in IMS-LD) and [Knowledge Resources]. A unit of learning depicts the flow of activities that a learner needs to follow in order to reach certain learning goals, giving certain prerequisites (Please note that in the remainder of this chapter, we simply use the term ‘learning activities’ to denote both units that are and are not modelled in *IMS-LD*). In our use case,

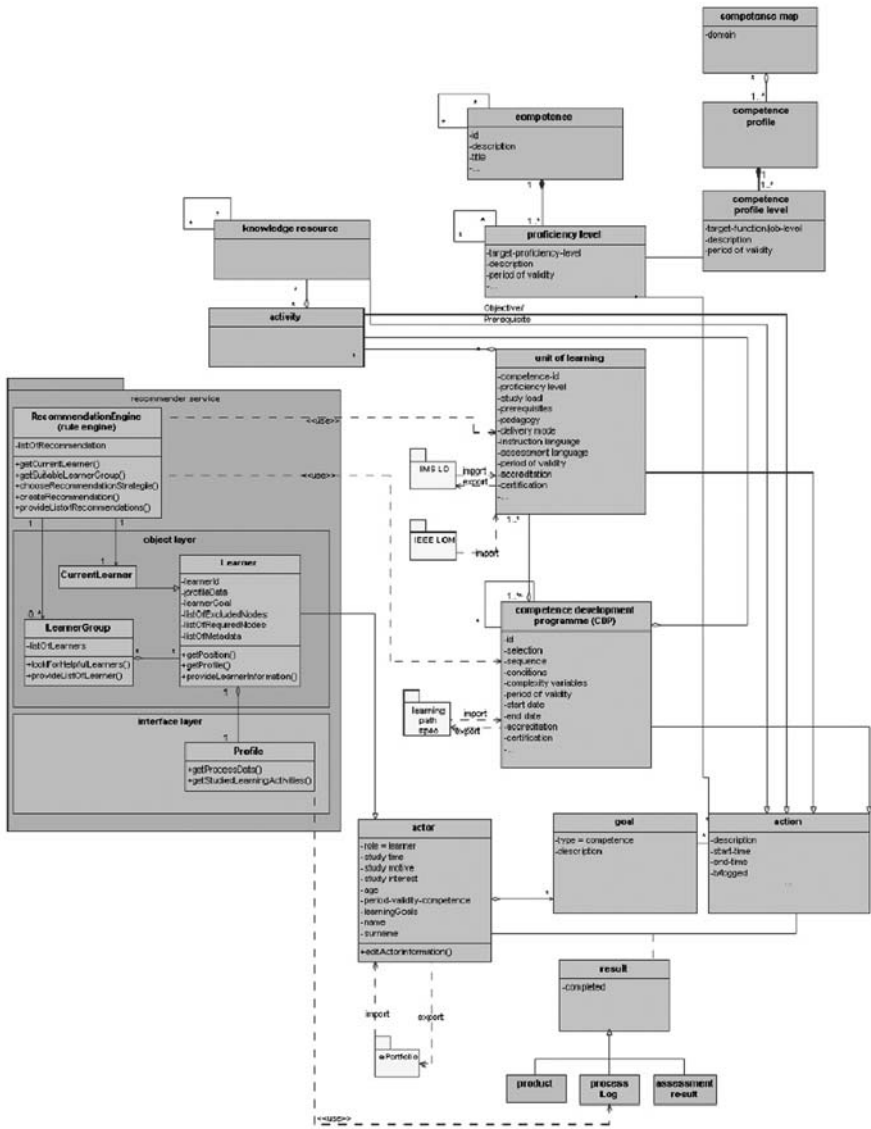


Fig. 6.2 Model for navigation services in Learning Networks

the learning goal is the completion of all 17 courses (equalling a total of 240 study hours), which all are on the same competence level, leading to the acquisition of the next (first) psychology competence level. Mastery of all courses, or goal attainment, will be assessed by a formal exam on location. Completion of each course is assessed by informal (facultative) progress tests provided with each course in the distributed Learning Network.

A learning path is an ordered set of actions, activities and/or units of learning that can be followed in order to acquire a competence with a certain proficiency level. Learning paths can include selections or sequences of ‘learning activities’, as well as conditions for their composition; attributes of a learning path can be modelled according to a *learning path specification*.

The core of the navigation services model is constituted by the [Recommender Service], where relevant information (from the classes just described) will be updated, integrated and matched to create personalised recommendations. In our use case learners store information about their needs and preferences (i.e., about available time to study, motivation to study, and psychological subdomain of interest) in a user profile (in Moodle), that can be adjusted during study. The recommender system we implemented took into account this user profile information and the completion of courses when advising the individual which course to study next, considering both individual data and collective data from the learner group of students with similar user profiles (other similar students that successfully completed X went on successfully completing Y). Depending on the learning situation (e.g., formal versus informal learning), a ‘recommendation strategy’ can be selected (e.g., more or less formalised information) by a [Recommendation Engine]. This engine takes into account learner positioning and profile information for each [Current Learner] (*information-based*) and/or for each [Learner Group] (*social-based*). In our use case this engine determined for each case whether sufficient data were available about the collective behaviour of the learner group. If this was not the case (cold start problem), the recommender fell back on information stored in the individual user profile.

The Learner Group class calculates and generates data about similar learners, in order to provide the engine with necessary matching information about the ‘peer group’. The [Learner] class gathers all required data about learners (actors with this role) from an interface layer to a *positioning service* (left out of scope in this model, leaving the feedback from the [Position] class null) and the [Profile] information. Our use case was relatively simple in the sense that no prerequisite knowledge or competence was required to enter the Learning Network, nor were exemptions granted based on such prior knowledge or competence. All profile information will be continuously updated when learning activities are completed (using the [Process Log] information) or when the ePortfolio is changed. The learner data are either (subjectively) provided or adapted by the learner, or (objectively) collected.

6.3.2 Experimentation with the Model

First versions of personalised recommender systems (PRS) have been developed at our centre over the last two years, according to the hybrid approach and navigation model that was sketched above. We have both set up a number of simulation studies and implemented a personalised recommender system in a real-life Learning

Network. The simulation results show that the presence of PRS leads to more satisfied and faster graduation. Furthermore, hybrid PRS using rating provide a good alternative for ontology-based recommendations. When setting up a simulation, all elements had to be defined, instantiated and stored, like: learner profile, competence map, recommendation strategy, goal, characteristics of the Learning Network (number of learners, activities, subdomains), preference gap, and competence gap. Carrying out simulation studies (using Netlogo as a tool) has proven to be a cost-effective way of determining the requirements and most optimal parameter settings for effective PRS, since they reveal the consequences of the assumptions that went into the design of the system without having to deal with ethical issues and practical constraints of providing a suboptimal system to actual learners (Berlanga et al. 2007; Nadolski et al. 2008), see also Chap. 8.

A more extensive and empirical study was carried out with a concrete navigation service deployed in an actual Learning Network, actually implementing the exemplary use case sketched in previous section (Sect. 6.3.2). A top-down, ontology-based recommendation technique and a bottom-up, stereotype filtering technique were combined in a recommendation strategy that decided which technique was most suitable for the current learner at the current moment. For practical reasons this system was limited to fixed learner profile metadata in a simple ontology. The current implementation of the system is limited to a fixed set of 17 formal learning activities that constitutes an Introductory Psychology course from one provider. About 250 learners entered this Learning Network, of which half were provided with the personalised recommender system. If only information about the learner was available, stereotype filtering was personalised through attributes of the learner profile. If it was not possible to base the advice on all attributes, the recommendation strategy would disable attributes (in a certain order) and the advice would be based on less attributes. Results showed that students receiving PRS advice completed more learning activities during study, needed less time to complete activities, created more personalised learning paths, and suffered lower drop-out rates (Drachler et al. 2008), see also Chap. 7.

6.4 Conclusion

This chapter started by arguing that professionals will need personalised navigation support in distributed Learning Networks, when choosing from the offerings of various content providers. Navigation services will then be needed to help the learner find and select most adequate activities and peers to study with. Personalised Recommender Systems are proposed to enable optimised learning paths and effective study. We have provided the reader with an overview of the main navigation concepts, learning technologies and recommendation techniques involved, also as an introduction to the chapters that will follow in this navigation section. It was found that an hybrid approach, combining information-based and social-based recommendation techniques into recommendation strategies is indeed a promising solution to both combine the benefits and limit the drawbacks of both approaches separately.

We proposed a model of the elements and attributes involved in such navigation services. It appears possible to design PRS that do not rely on rather specific and intensive data provisioning and maintenance, but could work mainly on low maintenance data derived from the behaviour of the Learning Network. For Learning Networks, it could therefore be concluded that a hybrid and layered ‘metadata ecology’ may be ideal, where the slow layers (ontologies) provide stability and the fast layers (folksonomies) drive change. Semantic web tools and learning technology standards provide a solid semantic framework (infrastructure) where there is a need for more formal and explicit characteristics of learners and learning activities to be decided upon from the top down. Where there is a need for less formal and implicit characteristics, free tagging and folksonomies offer more flexible, adaptive user feedback mechanisms to follow informal learning trends, emerging by serendipity and from the bottom up. Applied in practice, both simulation studies and an experimental field study have showed that PRS designed according to this hybrid approach yield to more, faster and more satisfied graduation, and indeed increase the variety of learning paths by self-organisation.

We do acknowledge current limitations when focusing on formal learning contexts with a fixed set of learning activities offered by one single institution. In such a situation, learners’ freedom in composing individualised selections and sequences of learning activities could be restricted by (institutional) constraints and prerequisites. Besides, personalisation assumes there will be one most suitable learning path for each individual learner, whereas the personalisation process actually might be more of an optimisation process of self-guided behaviour. In informal Learning Networks, the initial process of actively exploring and finding learning goals and learning opportunities might be more valuable than actually walking the learning path. Self-organisation assumes effective learning to simply emerge as active explorers learn from each others’ behaviour, which includes walking dead end streets as well. Therefore, consecutive systems are planned to broaden and generalise our findings to Learning Networks with contributions from various providers that entail both formal and informal learning activities.

Consecutive studies also have to examine ways to retrieve learner information as effortlessly and unobtrusively as possible. We do not have to burden professional, centralised indexers when ontologies could emerge entirely locally (some parts will be shared, some not, which is not an issue). For instance, McCalla (2004), opposed to the semantic-based paradigm, proposes a pragmatics-based approach of tagging learning activities with learner information. Each time a learner interacts with a learning activity, the (current) learner model of that learner is attached to the learning activity. It depends on the learning situation which learner data will be mined for patterns of particular use. More refined reasoning and recommendations will be possible when more and more instances accumulate. In addition to these instances, more standardised metadata could be assigned by professional indexers whenever needed. This ecological approach allows pre-assigned metadata (from ontologies like IEEE-LOM) to be refined or changed based on inferences from end use (from folksonomies tagged by learners).

Our roadmap for future research aims to combine collaborative filtering with both (a) information that can emerge bottom up by using rating or free tagging (folksonomies) of learning activities; and with (b) information that can be distilled from more formalised descriptions (specifications) of competences, learner profiles and learning paths. Future research is being planned to further establish the added value of personalising recommendations in terms of increased progress (effectiveness), decreased study time (efficiency) and experiences of users (satisfaction). Where possible, such information could be automatically added when learners interact with learning activities. When more specific information about learners and learning activities is required, metadata should be pre-tagged according to some ontology of attributes. In order to collect required information, learners will have to be encouraged to (automatically) update their ePortfolio or give ratings to (attributes of) the content. We will also need to establish the extent to which learners are willing to provide the necessary information about their personal needs and preferences. When they would not volunteer to do so but when considered necessary, the potential of incentive mechanisms (Hummel et al. 2005) could be considered to stimulate updating, tagging and rating.

References

- Agarwal, N., Haque, E., Liu, H., Parsons, L.: A subspace clustering framework for research group collaboration. *International Journal of Information Technology and Web Engineering* **1** (1), 35–58 (2006)
- Balabanovic, M., Shoham, Y.: Fab: Content-based, collaborative recommendation. *Communication of the ACM* **40** (3), 66–72 (1997)
- Berlanga, A.J., Van den Berg, B., Nadolski, R., Drachler, H., Hummel, H.G.K., Koper, R.: Towards a model for navigation support in Learning Networks. *E-Portfolio Conference* (Maastricht, The Netherlands, October 18–20, 2007)
- Berlanga, A., Sloep, P., Brouns, F., Bitter, M., Koper, R.: Towards a TENCompetence e-Portfolio. *TENCompetence Open Workshop* (Madrid, Spain, April 10–11, 2008)
- Bocchi, J., Eastman, J.K., Swift, C.O.: Retaining the online learner: profile of students in an online MBA program and implications for teaching them. *Journal of Education for Business* **79** (4), 245–253 (2004)
- Brand-Gruwel, S., Gerjets, P.: Instructional support for enhancing students' information problem solving ability [special issue]. *Computers in Human Behavior* **24** (3), 615–622 (2008)
- Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN-Systems* **30** (1–7), 107–117 (1998)
- Brusilovsky, P.: Adaptive hypermedia: user modelling and user-adapted interaction. In A. Kobsa (Ed.) *Ten Year Anniversary Issue Adaptive Hypermedia* **11** (1/2), pp 87–110 (2001)
- Burke, R.: Hybrid Web Recommender Systems. In P. Brusilovsky, A., Kobsa, W. Nejdl (Eds.). *The Adaptive Web: Methods and Strategies of Web Personalisation*, pp 377–408 (2007)
- CDM: A specification of Course Description Metadata. Proposal 20, October 2004. [Retrieved March 22, 2006 from: <http://cdm.utdanning.no/cdm/cdm-2.0.1/doc/courseDesc201004.pdf>] (2004)
- Cheetham, G., Chivers, G.: *Professions, Competence and Informal Learning* (Cheltenham: Edward Elgar Publishing 2005)
- Drachler, H., Hummel, H.G.K., Koper, R.: Effects of the ISIS Recommender System for navigation support in self-organised Learning Networks. *EduMedia Conference, Spe-*

- cial track on Technology Supported Self-Organised Learning (TSSOL)* (Salzburg, Austria, June 2–3, 2008)
- Guarino, N., Giaretta, P.: Ontologies and knowledge bases: towards a terminological clarification. In N. Mars (Ed.). *Towards Very Large Knowledge Bases* (Amsterdam: IOS Press 1995) pp 25–32
- Herlocker, J.L., Konstan, J.A., Terveen, L., Riedl, J.T.: Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*. **22** (1), 5–53 (2004)
- Hölscher, C., Strube, G.: Web search behavior of Internet experts and newbies. *Computer Networks* **33**, 227–346 (2000)
- HR-XML: *HR-XML Consortium: Competency Types* [Retrieved March 22, 2006 from: http://ns.hr-xml.org/2_4/HR-XML-2_4/CPO/CompetencyTypes.html] (2006)
- Hummel, H.G.K., Burgos, D., Tattersall, C., Brouns, F., Kurvers, H., Koper, R.: Encouraging contributions in learning networks using incentive mechanisms. *Journal of Computer Assisted Learning* **21** (5), 355–365 (2005)
- IEEE-RCD: *IEEE Standard for Learning Technology – Reusable Competency Definition* [Retrieved March 22, 2006 from: <http://ieeeltsc.org/wg20Comp>] (2006)
- IMS-ePortfolio: *IMS ePortfolio Specification* [Retrieved March 22, 2006 from: <http://www.imsglobal.org/ep/index.html>] (2004)
- IMS-LIP: *IMS Learner Information Package Specification* [Retrieved March 22, 2006 from: <http://www.imsglobal.org/profiles/index.html>] (2001)
- IMS-LD: *IMS Learning Design Specification* [Retrieved March 22, 2006 from: <http://www.imsglobal.org/learningdesign/index.html>] (2003)
- IMS-RDCEO: *IMS Reusable Definition of Competency or Educational Objective Specification* [Retrieved March 22, 2006 from: <http://www.imsglobal.org/competencies/index.html>] (2002)
- Janssen, J., Van den Berg, B., Tattersall, C., Hummel, H.G.K., Koper, R.: Navigational support in lifelong learning: enhancing effectiveness through indirect social navigation. *Interactive Learning Environments* **15** (2), 127–136 (2007)
- Janssen, J., Berlanga, A.J., Vogten, H., Koper, R.: Towards a learning path specification. *International Journal of Continuing Engineering Education and Lifelong Learning (IJCEEL)* **18** (1), 77–97 (2008)
- Koper, R.: Increasing learner retention in a simulated learning network using indirect social interaction. *Journal of Artificial Societies and Social Simulation* **8** (2) (2005)
- Koper, R., Sloep, P.B.: *Learning Networks: Connecting People, Organizations, Autonomous Agents and Learning Resources to Establish the Emergence of Lifelong Learning* [Report] Heerlen, The Netherlands: Open University of the Netherlands. [Available at www.learningnetworks.org] (2003)
- Lazonder, A.W., Biemans, H.J.A., Wopereis, I.G.J.H.: Differences between novice and experienced users in searching information on the World Wide Web. *Journal of American Society for Information Science* **51**, 576–581 (2000)
- Li, Q., Myaeng, S.H., Kim, B.M.: A probabilistic model for music recommendation considering audio features. *Information Retrieval Technology, Lecture Notes in Computer Science (LNCS)* **3689**, 72–83, Springer (2005)
- McCalla, G.: The ecological approach to the design of e-learning environments. *Journal of Interactive Media in Education* **7**, 1–23 (2004)
- Morville, P.: *Ambient findability* (Sebastopol, CA: O’Reilly 2005)
- Nadolski, R., Van den Berg, B., Berlanga, A., Drachler, H., Hummel, H., Koper, R., Sloep, P.: Simulating light-weight personalised recommender systems in learning networks: A case for pedagogy-oriented and Rating-based hybrid recommendation strategies. *Journal of Artificial Societies and Social Simulation (JASSS)*, **12**(1), 4 (2009).
- Ng, A., Hatala, M., Gasevic, D.: Ontology-based approach to learning objective formalization [Retrieved November 20, 2005 from: http://www.sfu.ca/lkng/fina-call_for_chapter.pdf] (2005)

- Oberle, D., Berendt, B., Hotho, A., Gonzalez, J.: Conceptual user tracking. In J.S.E.M. Ruiz, P.S. Szczepaniak (Eds.). *Web Intelligence. Proceedings of the First International Atlantic Web Intelligence Conference (AWIC 1003)* (Berlin: Springer 2003)
- Prins, F.J., Nadolski, R.J., Berlanga, A.J., Drachsler, H., Hummel, H.G.K., Koper, R.: Competence description for personal recommendations: The importance of identifying the complexity of learning and performance situations. *Journal of Educational Technology & Society*, **11**(3), 141–152 (2008).
- Rovai, A.P.: In search of higher persistence rates in distance education online programs. *Internet and Higher Education* **6**, 1–16 (2003)
- Schafer, J.B., Konstan, J., Riedl, J.: Recommender systems in e-commerce. *Proceedings of 1st ACM Conference on Electronic Commerce* (pp. 158–166), Denver, Colorado: Association of Computing Machinery (1999).
- Sicilia, M.-A.: Ontology-based competency management: infrastructures for the knowledge intensive learning organization. In M. Lytras, A. Naeve (Eds.): *Intelligent Learning Infrastructures for Knowledge Intensive Organisations: A Semantic Web Perspective*, pp 302–324 (Hershey, PA: IDEA Group 2005)
- Tattersall, C., Janssen, J., Van den Berg, B., Hummel, H.G.K., Koper, R.: Using IMS learning design to model curricula. *Interactive Learning Environments* **15** (2), 181–189 (2007)
- Van Setten, M.: *Supporting People in Finding Information: Hybrid Recommender Systems and Goal-Based Structuring* (Telematics Institute Fundamental Research Series No. 016, Enschede, The Netherlands 2005)
- XCRI: *eXchanging Course-Related Information* [Retrieved March 22, 2006 from: <http://www.elframework.org/projects/xcri>] (2006)

In this chapter we will describe the decisions developers have to make if they want to set up an experimental study to evaluate the effects of recommender systems for Learning Networks. Common tools for these kinds of feedback are recommender systems that support users in finding their way through the possibilities on the WWW. Many online companies like *amazon.com*, *netflix.com*, *drugstore.com*, or *ebay.com* (Linden et al. 2003; Schafer et al. 1999) are using a recommender system to direct the attention of their costumers to other products in their collection. The general purpose of recommender systems is to pre-select information a user might be interested in Adomavicius and Tuzhilin (2005).

Recommender systems can exist of various combinations of differently designed algorithms. A good overview about current recommender system technologies can be found in Burke (2002) and Herlocker et al. (2000). A detailed overview about especially recommender systems in the learning domain can be found (Drachsler et al. 2009b Nadolski et al. 2009).

This chapter offers guidelines to set up experiments for the evaluation of recommender systems in Learning Networks. It is based on our experiences with the ISIS experimentation that we conducted together with the Psychology department of the Open University of the Netherlands. The experiment focused on supporting learners in their course selection by providing personalised recommendations. In this chapter we focus on the methodology and technical decision that have to be taken; detailed experimental results and further information of the ISIS experiment can be found in Drachsler et al. (2009b).

In order to design an experiment for recommender systems in Learning Networks several things have to be considered. Firstly, the experimental designers have to be aware of the underlying concept of professional development that inspires the whole experiment. Secondly, they have to be aware of the Learning Networks concepts. Thirdly, the designers are expected to have at least basic knowledge about recommender system technologies. Finally, standardised methods of experimental design are required in order to run a valid experiment. In the following section of the chapter we shortly introduce the reader to those requirements.

First, for a proper experimental design researchers have to decide which hypotheses should be tested and which variables support those hypotheses. Most of the times in Technology Enhanced Learning (TEL) research, we want to observe if learners perform more efficient, more effective, are more satisfied, or if the instrument used decreased the drop-out rate of learners. In the special case of Learning Networks we also have to consider aspects form Social Networks Analysis in order to analysis how the network benefits from the contributions of their learners.

Second, running real life experiments with recommender systems requires a specific kind of statistic analysis. This analysis is based on measurements on a regular basis over a fixed time period. It enables the researches to monitor the effects of their instrument (recommender system) during the runtime and at the end of the experimentation.

Third, experimental designers have to make a decision which techniques should be used to present the learning activities to the participants of the experiment. Most of the time, a common virtual learning environment (VLE) will be selected. There

are many Open Source solutions available like *Drupal* or *Moodle* to set up a Learning Network. The experiment can also rely on an alternative in-house solution that is already successfully applied in an institution. Especially for recommender systems, researchers have to make a decision if they want to build their own recommender system or apply already existing recommender system plug-ins or frameworks in a VLE.

In the second section of this chapter, we will now describe an experimental design for a recommender system in a Learning Network. Section three explains details about the statistical analysis of this ISIS experiment. Section four will discuss the selection of suitable techniques. Finally, the last section offers ideas for future research regarding recommender systems in Learning Networks.

7.2 Experimental Design

In the recommender system research, most of the time offline experiments are done with several data sets with specific characteristics (the MovieLens dataset, the Book-Crossing data sets, or the EachMovie dataset) before preparing an experiment with real users (Goldberg et al. 2001; O'Sullivan et al. 2002; Sarwar et al. 2002). This is also because classic recommender system research has its focus on the optimisation or invention of more efficient algorithms for certain recommendation problems. These data sets are used as a common standard or benchmark to evaluate new kinds of recommendation algorithms. Furthermore, machine-learning research only evaluates recommendation algorithms based on common technical measures like accuracy, coverage, and performance in terms of execution time (Adomavicius and Tuzhilin 2005; Burke 2002; Herlocker et al. 2004). Accuracy measures how close the predicted ranking of items for a user differs from the user's true ranking of preference. Coverage measures the percentage of items for which a recommender system is capable of making predictions. Performance observes if a recommender system is able to provide a recommendation in a reasonable time frame.

Research on recommender systems in Learning Networks is also in need of these technical measures, but in the first place we have to improve the learning process with the selected technology. We have to deal with information about learners and learning activities and combine different levels of complexity for the different learning situations the learner may be involved in. The main recommendation goal for recommender system in Learning Networks is to provide learners with suitable learning activities in order to support their professional competence development. Therefore, recommender systems in Learning Networks have to consider relevant pedagogical rules describing pedagogy-oriented relations between learners' characteristics and LA-characteristics. For example: from Vygotsky's 'zone of proximal development' follows the pedagogical rule 'recommended learning activities should have a level a little bit above learners' current competence level' (Vygotsky 1978). Thus, recommender systems in Learning Networks have to take into account competence levels in order to suggest an appropriate learning activity. Further differences between recommendation in the e-commerce domain and the learning domain can be found in Drachsler et al. (2009a).

Currently, we do not have any standardised data sets for offline experiments publicly available. Further, it is not appropriate to focus only on technical measures for recommender systems in Learning Networks without considering the actual needs and characteristics of professionals. Thus, further evaluation procedures that are complementary to technical evaluation approaches are needed.

In the following we split this section into two subsections. The first subsection (Sect. 7.2.1) explains general requirements to evaluate recommender system in Learning Networks. The second subsection (Sect. 7.2.2) describes the experimental setup of the ISIS experiment in detail.

7.2.1 An Evaluation Framework for Recommender Systems in Learning Networks

A pedagogy driven recommender system for Learning Networks that takes into account learner characteristics and specific learning demands also should be evaluated by multiple evaluation criteria. To evaluate the influence of the recommender system we need a mixture of educational, technical and network measures. We advise you to mix technical evaluation criteria with educational research measures and network measures (Drachsler et al. 2009a) in a recommendation framework. Therefore, we suggest the following for the analysis of the suitability of recommender system in Learning Networks.

Classic educational research is most of the time evaluated base on the outcomes of the learning process of the learner (Thorpe 1988). The aim is to develop the competences of the learner on cognitive or motor level. Therefore, commonly used measures for valid evaluations are effectiveness, efficiency, satisfaction, and the drop-out rate because of two reasons. First, these criteria are used to evaluate for instance universities regarding their outcomes, and second they can be efficiently operationalised. For example, effectiveness is a measure of the total amount of completed, visited, or studied learning activities during a learning phase. Efficiency indicates the time that learners needed to reach their learning goal. It is related to the effectiveness variable through counting the actually study time. Satisfaction reflects the individual satisfaction of the learners with the given recommendations. Satisfaction is close to the motivation of a learner and therefore a rather important measure for learning. Finally, the drop-out rate mirrors the numbers of learners that dropped out during the learning phase. In educational research the drop-out rate is a

Table 7.1 Suitable measurements and their corresponding parameters building up an evaluation framework for recommender systems in Learning Networks

Measurements	Parameters
Technical measures	Accuracy Coverage Performance
Educational measures	Effectiveness Efficiency Satisfaction Drop-out rate
Social network measures	Variety Centrality Closeness Cohesion

very important measure because one aim is to graduate as many learners as possible during a learning phase.

The Social Network Analysis (SNA) measures are needed to estimate the benefit coming from the contributions of the learners for the network as a whole (Wasserman and Faust 1999). These are more specific measures that are mainly related to informal Learning Networks. SNA give us various insights into the different roles learners own in a Learning Network. SNA measures like variety, centrality, closeness and cohesion. Variety measures the level of emergence in a Learning Network through the combination of individual learning paths to the most successful learning routes. Centrality is an indicator for the connectivity of a learner in a Learning Network. It counts the number of ties to other learners in the network. Closeness measures the degree a learner is close to all other learners in a network. It represents the ability to access information direct or indirect through the connection to other network members. Cohesion indicates how strong learners are directly connected to each other by cohesive bonds. Peer groups of learners can be identified if every learner is directly tied to every other learner in the Learning Network.

These evaluation criteria can be conflicting. For instance, learners with many rated learning activities get a central role in a Learning Network from the SNA perspective. They get many direct ties to other learners through the huge amount of rated learning activities. From an SNA perspective these learners are beneficial for the Learning Network because they contribute heavily to it. But from the educational research perspective the same group of learners may be less important because their educational measures are quite poor. It might be that they needed much more study time (efficiency) or complete less learning activities successfully (effectiveness) compared to others learners in a Learning Network (LN). Therefore, further research regarding the evaluation of recommender systems for their support for learners in LNs is needed.

7.2.2 An Exemplary Experimental Setup to Evaluate a Recommender System in a Learning Network

To evaluate a recommender system in a Learning Networks we conducted, together with the Psychology faculty of the Open University of the Netherlands, the ISIS experiment. In ISIS (Individualised Support In Sequencing) the learners were able to study learning activity in any order instead of following the learning activities in a fixed order. The experiment focused on supporting learners in their course selection through personalised recommendation by a recommender system. The recommender system supported them with recommendations based on their learner profile and the behaviour of learners that were similar to them. We called that approach personalised navigation support and were especially interested in the learning outcomes of the learners and less in measures like algorithm performance of the machine-learning field. Thus, we selected effectiveness, efficiency, variety and satisfaction as evaluation criteria from the evaluation framework.

The following hypotheses were tested in the ISIS experiment, where the control group was provided with the Moodle virtual learning environment and a text book; whereas the experimental group was additionally provided with a recommender system that recommended best next learning activity based on successful choices of other learners with similar profiles.

The experimental group will be able to complete more learning activities than the control group (effectiveness). The proportion of completed learning activities is bigger in the experimental group compared to the control group.

The experimental group will complete (the same amount of) learning activities in less time, because alignment of learner and learning activity characteristics will increase the efficiency of the learning process (efficiency).

The experimental group has a broader variety of learning paths than the control group because the recommender system supports more personalised navigation (variety).

The experimental group will be satisfied with the navigational support of the recommender system (satisfaction).

It is always challenging to design an experiment corresponding to real life conditions because conditions are never the same like in a laboratory. However, the experimental design has to be strict as possible. In our example we adapted a formal course of the Psychology faculty of the Open University of the Netherlands to certain characteristics of professionals in Learning Networks. Consequently, we used the learning activities designed by domain experts and integrated them into a condition which was comparable to a Learning Network.

In the ISIS experiment we focused on the delivery of learning activities to professionals. We neglected the learning activity creation by learners and focused purely on learner support through recommender systems. In order to draw conclusions to professional development networks we especially addressed professional development characteristics like self-responsibility and its support through recommender systems. Therefore, we neglected the formal university conditions and constraints to design the experiment as similar as possible to the conditions of professionals in Learning Networks. Both groups got a maximum of freedom for their studies; in principle they were able to study the course over years. We informed both groups that they do not have to follow the learning activities in a certain order or pace. Further, the students could register for a final exam whenever they wanted, even without completing any of the online available multiple-choice tests for self-assessment.

Detailed results of the ISIS experiment that acts here as an example can be found in Drachsler et al. (2009b). The experiment examined the effects of the navigation support on the completion of learning activities measured (effectiveness), needed time to complete them (efficiency), satisfaction with the system (satisfaction), and the variety of learning paths (variety). The recommender system positively influenced all measures with having significant effects on efficiency, variety, and satisfaction on a four month run time.

Participants. In order to run experiments with recommender systems in Learning Networks the experimental designers should get as many participants as possible, because there is always a drop-out rate on various levels of participation. Thus, the group of participants that can be used for statistical analysis is getting smaller than the initial number of subscriptions.

In our example a total of 244 participants subscribed to the ISIS experiment. All participants were distance learners who studied the learning material on their own. Both the experimental and control group contained an equal amount of learners (122 learners per group) because the learners were randomly allocated, see Fig. 7.1. Twenty-four participants (19.7%) in the experimental group and 30 participants (24.5%) in the control group never logged into the Moodle environment. This group of non-starters was not included in our analyses. This leaves a group of 190 learners who did enter the Moodle environment; 98 in the experimental and 92 in the control group.

The group of actual starters had to be further differentiated into active and passive learners, because not all of the learners actually used or made progress in the Moodle environment. From the 98 participants in the experimental group 72 learners completed learning activities; from the control group 60 learners completed learning activities. Thus, in total a group of 132 were active learners during the experiment. We used this total amount of active learners to analyze hypothesis 1 (Effectiveness), hypothesis 2 (Efficiency), and hypothesis 3 (Variety).

The participants could voluntarily register for the new version of the course, and were informed that they were taking part in an experiment with a new learning environment. They were not informed that only half of the students would receive additional navigation support.

The conditions of the experiment allowed learners to start their studies whenever they want to. As a consequence not all students started at the same time; some of them started later and we got a *dynamic starting point* of students that have to be specially treated in the statistic analysis.

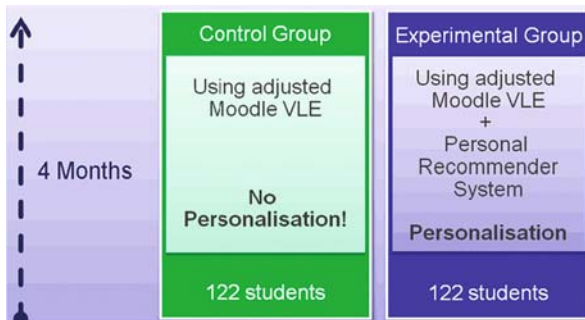


Fig. 7.1 Experimental design of the ISIS experiment

7.3 Statistical Analysis

To evaluate the effects for the experiment according to our hypotheses we applied a mix of different analysis procedures.

Effectives and *efficiency* measures were monitored every two weeks during the experimental runtime with a *repeated measurement design*. The repeated measurement design is part of the *generalised linear model* (GLM) a flexible generalisation of ordinary least squares regression. The GLM is commonly used in applied and social research. It is the foundation for the t-test and the Analysis of Variance (ANOVA).

For the evaluation of the *variety* of learning paths we developed a visualisation tool based on the multi-agent environment Netlogo. The tool shows an overlay of all learning paths within a group of learners. Thus, you can easily recognise their variance in the learning paths.

Satisfaction was measured through an online questionnaire and further analyzed with descriptive statistics. Therefore, we used an Open Source questionnaire tool called UCCASS (<http://www.bigredspark.com/survey.html>). In the following sections we introduce the different analysis techniques for the ISIS experiment.

7.3.1 Analysis of Effectiveness and Efficiency

In order to deal with a selection problem in our experiment we defined a goal attainment of 5 completed learning activities out of 17 in total. Our aim was to support as much learners as possible to complete these 5 learning activities as fast as possible. To measure the effectiveness and efficiency of the recommender system learners were taken into account that applied to the following condition; completed more than 5 learning activities, or successfully completed the final exam, or were still studying at the measure point. This condition leaves a number of 101 students at the end of the experiment ($n = 52$ in the experimental group and $n = 49$ in the control group). Regarding the individual dynamic starting points of the students the recorded measure in Table 7.1 contained 0 values in case students started later. In order to run a MANOVA analysis (Keselman et al. 1998) all individual starting points of the students were moved in one 'starting' column through deleting the 0 values. Therefore, Table 7.1 was transformed into a study progress table (see Table 7.2). Table 7.2 differs from Table 7.1 through moving the individual starting points into one 'starting' column (first column), and duplicating the study results towards the end of the table if the students complied to the above mentioned condition. To test hypothesis 1 and 2, we analysed the measures taken using SPSS version 12. To avoid inflated Type I error due to multiple tests, a priori tests of specific contrast scores were used.

The effectiveness and efficiency was analyzed by means of linear and quadratic trend analysis. To test hypothesis 1 and 2, we analysed the measures taken using SPSS version 12. To avoid inflated Type I error due to multiple tests, a priori tests of specific contrast scores were used. The effectiveness and efficiency was Averaged completion scores and averaged completion time during the two experi-

Table 7.2 Example table of biweekly recorded measures

Learner	Biweekly measure points						
	Oct	Oct 2	Nov	Nov 2	Dec	Dec 2	Jan
1	1	2	4	7	7	7	8
2	0	0	0	1	3	5	9
3	0	0	0	0	0	1	1
<i>n</i>	1	2	3	4	4	4	4

This table represents the not yet transformed recorded measures of the biweekly measure points. The 0 values are related to the individual starting point of the participants. The numbers show the amount of learning activities they completed successfully at the specific measure point.

Table 7.3 Example table of prepared biweekly measures for MANOVA analysis

Learner	Study progress per learner per measure point						
	1	2	3	4	5	6	7
1	1	2	4	7	7	7	8
2	1	3	5	9	9	9	9
3	1	1					
<i>n</i>	1	2	3	4	4	4	4

This table shows the actual study progress of all active learners. Therefore, all 0 values from Table 7.1 are deleted and the individual starting points were moved into one 'starting' column (first column). The MANOVA analysis in SPSS requires equally distributed values for each participant. If the learners completed more than 5 learning activities or they completed the final exam and not for each column a value was available their final study result was duplicated towards the final measure point (e.g. Learner 2). Learners that completed less than 5 learning activities were only taken into account when they still studied at the final measure point (e.g. Learner 4). Learners like learner 3 were not taken into account because they did not complete more than 5 learning activities and were not studying at the final measure point.

mental periods were transformed into linear and quadratic trend contrast scores by means of computation of orthogonal polynomials. We applied multivariate analysis of variance (MANOVA) for repeated measures on these a priori chosen contrast scores with Group as between subjects factor and Time as within subjects factor. A significant interaction of contrast scores with Group was followed by testing of simple contrast effects. Due to the a priori character of these tests, they were performed with the conventional Type I error of 0.05 (Tabachnick and Fidell 2001).

7.3.2 Analysis of Variety of Learning Paths

To test hypothesis 3, the variety of learning paths, we analyzed the behaviour of the learners with a Graph Theory approach (Gross and Yellen 2006). Therefore, we modelled the Learning Network in Netlogo 4 (Tisue and Wilensky 2004), and observed the completion of learning activities by the learners. Analysis software and

example data set can be downloaded (<http://hdl.handle.net/1820/1493>). If a learner completed for instance first learning activity 1 and second learning activity 7 it was counted as traffic between learning activity 1 and learning activity 7. A line was drawn between both learning activities in the graph when the traffic became larger than 3.




Fig. 7.2 Example picture of the variety of the learning paths. The standard curriculum order is indicated through numbers. *Arrows* show the learning paths of the learners in a group

If the learning path was used even more frequently, the traffic line got thicker and changed its colour. Consequently, the thickest path was used most often and the thinnest path was used only three times.

7.3.3 Analysis of Satisfaction with the Recommender System

To test hypothesis 4, the general satisfaction of the recommender system, we conducted an online recall questionnaire. The questionnaire was sent to all participants in both groups at the end of the experiment.

The Open Source UCCASS system makes online questionnaire an easy procedure. The system is also based on PHP and MySQL and therefore adjustable for

How often did you follow the advice of the recommender system?




Fig. 7.3 Screenshot of the result view of the UCCAS online questionnaire system

certain wishes. It offers the possibility to load all participants into the MySQL database and to submit an invitation to every participant via e-mail. Further, any common question design is available and the results of the questionnaire can be filtered on different levels. The questionnaire results can easily exported from the database integrate in statistic programs like SPSS.

7.4 Suitable Recommendation Systems and Techniques

Depending on your resources and on the purpose of your experiment you have the choice between already existing recommender system plug-ins, programmable frameworks, or toolkits with additional functionality. There are also plenty of scientific publications regarding recommender system techniques which can be used to program own recommender systems (Adomavicius and Tuzhilin 2005; Burke 2002; Herlocker et al. 2004).

In the following section we will discuss various recommendation plug-ins, frameworks and a toolkit that can help to set up a recommender system environment for an experiment in Learning Networks research. Detailed information about recommender system techniques and how they can be adapted to the specific purposes for Learning Networks can be found in Drachsler et al. (2009b).

7.4.1 Available Recommender Systems

Currently, several recommender systems are available on various complexity levels. Some of them are available as plug-in for VLEs and websites and others are frameworks that have to be instantiated. Instantiations require programming effort but using a framework is still easier than creating an own recommender system from the beginning. A major advantage of the frameworks is that experimental researches

Table 7.4 Available recommender systems

Software type	Recommender systems
Plug-ins	Recommendation module for Drupal Vogoo
Frameworks	CoFE Taste Duine
Toolkit	Scout portal

can be sure to use the most efficient and effective recommendation algorithm from the machine learning field without being confronted with the mathematical calculations behind the algorithms. Instead of that the researchers have to feed the system with learning activities and learner profile information. The following systems are available.

On the plug-in side there are two suitable systems available, a *Content Recommendation Engine for Drupal* and the *Vogoo* recommender system. Both are based on PHP code and therefore easily to integrate into PHP based VLE like Moodle or Drupal.

The easiest way to integrate a recommender system into a VLE is the recommendation module for Drupal (<http://drupal.org/node/920>). It is limited to user-based collaborative filtering only. The module recommends interesting nodes, according to personal tastes of a user compared with other users in the system. Thus, users have to rate a couple of nodes (as ‘Not Recommended’, ‘Recommended’, or ‘Highly Recommended’) in order to get recommendations.

Another possibility is the *Vogoo PHP Lib* (<http://www.vogoo-api.com/>) a free PHP library licensed under the terms of the GNU GPL v2. The *Vogoo PHP Lib* has been designed with ease-of-use in mind. The team promises to add professional collaborative filtering functions to a website in minutes. *Vogoo PHP* includes two item-based and one user-based collaborative filtering technique and is therefore more flexible than the Drupal module. The *Vogoo* team also offers a commercial version called *Vogoo PHP Pro* as a proprietary version of *Vogoo PHP Lib*. This includes a highly optimised pre-computation engine for item-based collaborative filtering. Performance tests have shown an improvement of up to 20 times in execution speed for pre-computation scripts when compared to the GPL version.

On the framework side you can choose between three different recommender systems the *Taste*, the *CoFE*, and the *Duine* framework. An advantage of the recommender frameworks is the possibility to adapt the recommendation task to specific requirements of your experiment or your domain. This is not possible with the plug-ins, because they offer less flexibility for further development.

Taste (<http://taste.sourceforge.net/>) is a flexible collaborative filtering engine written in Java. It can be used as standalone application but it also can be used as external server, which exposes recommendation logic to your application via web services. The engine takes users preferences for items (‘tastes’) and returns estimated preferences for other items. *Taste* provides a rich set of components from which you can construct a customised recommender system from a selection of

algorithms. It addresses important recommender system issue like performance, scalability and flexibility to provide fast recommendations also for huge data sets.

A similar project is the CoFE (<http://eecs.oregonstate.edu/iis/CoFE/>) project developed by the Intelligent Information Systems research group of Oregon State University. CoFE is a free, Open Source server for the Java platform that anyone can use to set up a recommendation system. Features include individual items recommendations, top-N recommendations across all items, top-N recommendations based on one type of item. Recommendations are computed using a popular, well-tested nearest-neighbour algorithm (Pearson's algorithm).

Finally, the Duine (<http://sourceforge.net/projects/duine/>) framework allows users to develop own prediction engines for recommender systems. Duine is also Open Source and available for free. Duine contains a set of recommendation techniques, ways to combine these techniques into recommendation strategies, a profile manager, and it allows users to add their own recommender algorithm to the system. Duine already includes some further functionality like a component for the management of user profiles. The result of a Duine prediction engine is the retrieved set of information with added data about how interesting each piece of information is for the user.

In the category toolkits, the Scout Portal Toolkit is available which makes it possible to set up a whole content management system. It can also be used to set up a VLE such as the learning languages project (<http://www.learninglanguages.net/>) that makes advantage of it. It is one of the easiest and fastest ways to setup an experiment for Learning Networks including a recommender system. The Scout Portal Toolkit provides a number of features beside a recommender system. It also enables cross-field searching, resource annotations by users, intelligent user agents, and resource quality ratings by users. The recommender system uses item-based filtering technique, based on community ratings.

Most of the time one of the presented recommender systems is more suitable than other ones for certain research conditions. If researchers want to make a case study within the concept of Learning Networks and no programming capacity is available we suggest using the Scout Portal Toolkit. In this case, the setup of the Learning Network is rapidly done and it already contains a recommender system. Similar applies for the Voogoo and the Drupal plug-in. In both cases the experimental team has to add learning activities to a VLE and can additionally add a recommender system with minor programming knowledge. The Scout Portal Toolkit and the Drupal plug-in are both based on one recommendation technique only. The Voogoo plug-in offers already three different recommendation techniques but therefore it is also a bit more challenging regarding the implementation.

If the experimental designers have more specific research questions regarding recommender system in Learning Networks we suggest to use one of the recommender system frameworks. They allow much more adjustments of the systems to any experimental design and still hide complexity of recommender system algorithm. However, they definitely require more programming capacity and a deeper understanding of recommender system insights than the other solutions.

In case experimental designers decide to design a recommender system from the bottom onwards, they have the most freedom and possibilities for the development of a specific recommender system for a certain recommendation task. There are three overview articles available that are supportive for a selection of the most suitable recommendation technique (Adomavicius and Tuzhilin 2005; Burke 2002; Herlocker et al. 2004). For the ISIS experiment we decided to develop our own recommender system with particular aspects regarding professional development in distributed Learning Networks. We did so because we collaborated with the Psychology faculty at our institute that wanted to evaluate the *Moodle LMS* for their distance courses. In the joined ISIS project we supported them to set up and gain experience with Moodle. This way, we could rely on the learning material and the students as participants for our experiment. At the end of the ISIS project Psychology was satisfied with the research results and decided to use Moodle as LMS for all courses as well as to use the recommender system. Currently, they further develop the experimental prototype of the recommender system for support in other courses as well.

7.4.2 The Techniques We Used in the ISIS Experiment

For the ISIS experiment we decided to combine a domain ontology with a stereotype filtering technique. Recommender systems with a combined recommendation strategy provide more accurate recommendations when compared to single techniques recommender systems (Melville et al. 2002; Pazzani 1999; Soboro and Nicholas 2000). The ontology used personal information of the learner (e.g., interest) and compared that with the domain knowledge to recommend the most suitable learning activity. Stereotype filtering used profile attributes of the learners (e.g., interest, motivation, study time) to create learner groups and recommend learning activities preferred by similar learners.

The recommender system advises the next best learning activity to follow based on the interest of learners (ontology-based recommendation), and on the behaviour of the peers (stereotype filtering). If only information about the interest of a learner was available, then ontology-based recommendation technique was used, else the stereotype filtering technique was applied. The underlying recommendation strategy is presented in Fig. 7.4.

The use of the stereotype filtering was prioritised and the ontology approach was used mainly to cover the 'cold-start problem' (Herlocker et al. 2000) of the stereotype filtering technique. The stereotype filtering technique was personalised through attributes of the personal profile of the learners. If it was not possible to give any advice it disabled one of the personal attributes and tried to make a recommendation based on larger peer group with less common attributes.

Only in the case that the stereotype filtering was not able to provide any recommendation, the recommender system created ontology-based recommendations. The ontology visualised in Fig. 7.5 consists of two top domains (e.g., 'Environmen-

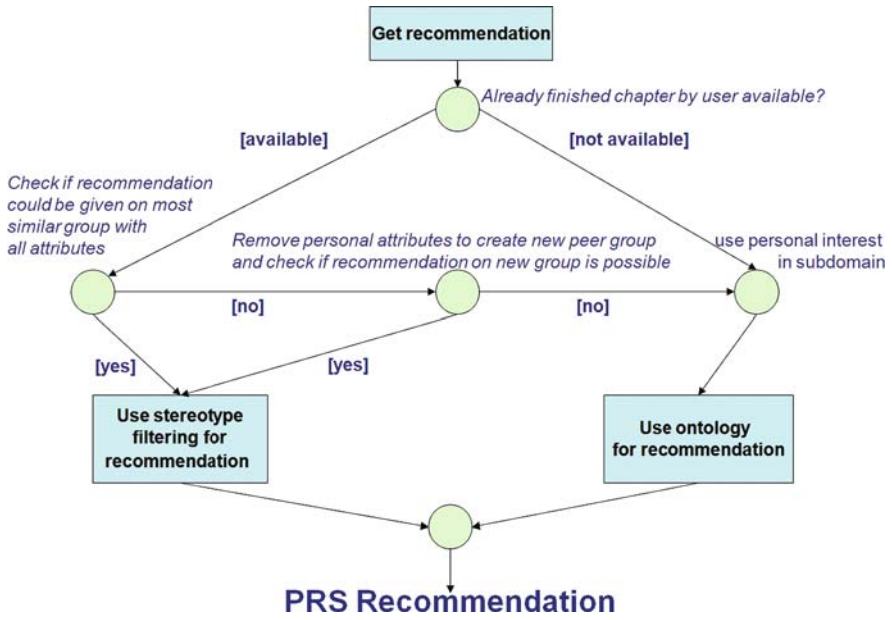


Fig. 7.4 Recommendation strategy for the implemented recommender system

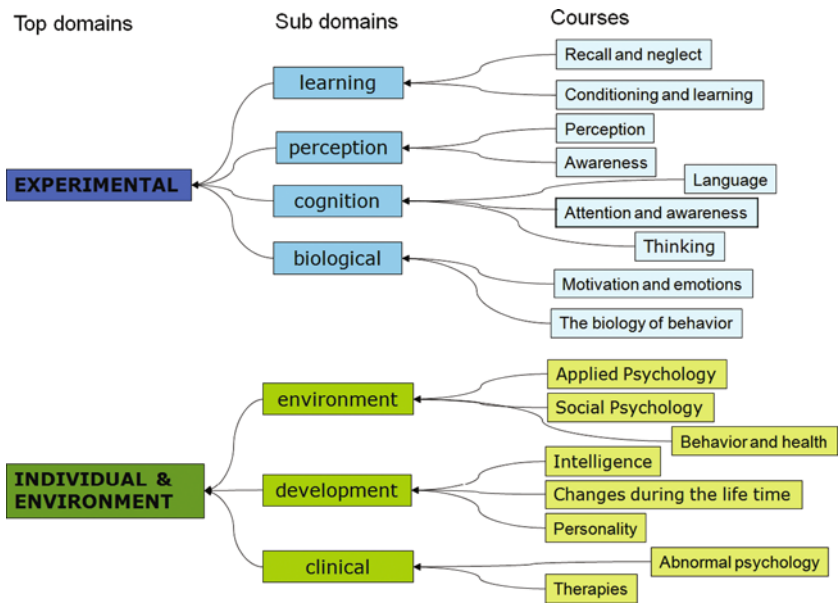


Fig. 7.5 Structure for ontology based recommendations

tal Psychology’) that contain several sub domains (e.g., ‘learning’), each containing two or three courses (or learning activity) (e.g., ‘recall and neglect’). The learners had to select a special interest (one of the sub domains of the ontology) in their profile. If the learners had chosen a sub domain (e.g., ‘clinical’), they received recommendations on courses located in that particular sub domain. If none of these courses had been completed by others so far, the recommender system randomly recommended one of them. If one course had already been completed by the learner the other course(s) was/were recommended. If all courses of the sub domain (e.g., ‘clinical’) were completed the ontology recommended a course that was part of the top domain ‘Environmental Psychology’.

7.4.3 The Virtual Learning Environment

We selected Moodle as VLE (Dougiamas 2007), because it is an Open Source solution written in the PHP programming language and therefore easily adaptable to our experimental needs. The learning activities and the recommender system were implemented into Moodle. Moodle was adjusted to the experimental setup, thus some functionality of Moodle was blurred out and other functionalities like a multiple-choice tool where additionally added. Figure 7.6 shows the overview screen of learning activities for a learner in the experimental group. The overview is divided into three columns. The right column shows the learning activities the learner still has to study. The middle column presents the courses the learner is already enrolled for. Finally, in the left column all completed courses are listed.


Overview of learning activities		
<p>You already completed: You have not completed any learning activity.</p>	<p>Activities you are enrolled into: Perception Personality Awareness Changes during the life time Therapies Language</p>	<p>You still need to complete: Behavior and health Thinking Social Psychology Conditioning and learning Abnormal psychology Recall and neglect Intelligence The biology of behavior Motivation and emotions Attention and awareness Applied Psychology</p>
 <p>Based on your study interest in "cognition" (mentioned in your personal profile), we suggest to further study the following learning activity:</p>		
Title of the suggested learning activity		Options
Thinking		description of the recommendation adjust profile

Fig. 7.6 Overview page of the experimental group with a recommendation

Below an explanation of the recommendation is given. In this screen, the recommender system has recommended ‘Thinking’ as next best course. Next to the recommendation there are additional options to get further information about the recommendation and to adjust the preferences set in the learner profile.

The Learning Network that was based on a Moodle adaptation contained 17 learning activities with an average study load of 12 hours. Completion of each learning activity was assessed by multiple-choice tests consisting of seven equally weighted questions. A score of 60% or more was considered as a successful completion of the learning activity. With the Moodle environment the learners received an Introduction to Psychology handbook that contained additional information to the 17 learning activities. All learning activities were separate entities in Moodle, setup according to the same didactical structure. The Moodle environment contained all further learning materials, including support and guidance, task assignments, progress tests, additional pictures and links, summaries, and other attractive learning tasks.

7.5 Conclusion

We have presented all the required tools and concepts that are needed to set up an experiment with recommender systems in Learning Networks for professional development. We have given an overview about a suitable experimental design and offered an example for that. Further, we introduced statistic methods and procedures to test hypotheses that can be based on a selection of variables from an evaluation framework. Finally, we discussed various available recommender system and suitable virtual learning environments to create a Learning Network. In this final section we want to give incentives for future research on the navigation support through recommender systems in Learning Networks.

Following experiments in this field can vary on four key elements: Changing the underlying recommendation algorithms, Adjusting the pedagogic context, Addressing a specific user group (older people, more technologically literate, higher educational achievement), and Using a different VLE or other educational services for the experiment.

These four key elements can be combined in various experimental settings. Based on the ISIS experiences we suggest to continue with variations on the second and fourth elements. We aim to apply the use of informal learning activities created by the professional to address the navigation problem in Learning Network on a higher level. Research in this area should make advantage of learning activities available in Web 2.0 services like *wikipedia.com*, *youtube.com* or *slideshare.com*. Future experiments in this area should use a mixture of formal and informal learning activities to simulate a Learning Network. In this case, it is hardly possible to apply a domain ontology because of the ‘open corpus problem’ (Brusilovsky and Henze 2007). The open corpus problem applies when an unlimited set of documents are given that can not be manually structured and indexed with domain concepts and metadata from a community. Thus, to prepare recommendations for informal learning activities dif-

ferent recommendation strategies have to be invented. Therefore, Open Educational Resources (OER) (Hylén 2006) are also a very interesting source for the data base of future experiments in Learning Networks. Experimental designers should consider mixing different kinds of these OER repositories and maybe additionally combining them with learning activities created by learners.

An unsolved issue is the measurement of accepted recommendations by the learner. The problem is the definition of an ‘accepted recommendation’. Did learners appreciate a recommendation when they navigated to a recommended learning activity? Or did learners accept a recommendation when they used the recommended learning activity more than 5 min? Anyway an objective measure is needed to indicate a successful recommendation for a learning activity. In e-commerce recommender system a recommendation was successful if a consumer finally bought a recommended product. In the case of professional development we have to measure at least that a learner is busy with a learning activity. This could be done with various indicators like ‘time spend on learning activity’, ‘click rate’, ‘repeated use of the learning activity’, and ‘added content to learning activity’ in an interaction model.

Finally, in the ISIS experiment we limited ourselves to show only the ‘best next learning activity’, based on our recommendation strategy to the learners. We did that for experimental reasons. It is also thinkable to select a different experimental design and offer sorted lists of recommendations. In the real life of professionals a list or a sequence with suitable recommendations might be more valuable than a single recommendation.

References

- Adomavicius, G., Tuzhilin, A.: Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* **17**(6), 734–749 (2005)
- Brockett, R.G., Hiemstra, R.: *Self-direction in Adult Learning: Perspectives on Theory, Research and Practice* (London: Routledge 1991)
- Brusilovsky, P., Henze, N.: Open corpus adaptive educational hypermedia. In P. Brusilovsky, A. Kobsa, W. Nejdl (Eds.): *The Adaptive Web: Methods and Strategies of Web Personalization* Lecture Notes in Computer Science ed., Vol. 4321, pp. 671–696 (Berlin, Heidelberg, New York: Springer 2007)
- Burke, R.: Hybrid recommender systems: survey and experiments. *User Modeling and User-Adapted Interaction* **12**(4), 331–370 (2002)
- Dougiamas, M.: *Moodle* [Retrieved 12 June 2007, from <http://moodle.org/>] (2007)
- Drachsler, H., Hummel, H., Koper, R.: Personal recommender systems for learners in lifelong learning: requirements, techniques and model. *International Journal of Learning Technology*, **3**(4), 404–423 (2008)
- Drachsler, H., Hummel, H., Koper, R.: Identifying the goal, user model and conditions of recommender systems for formal and informal learning. *Journal of Digital Information*, **10**(2), 4–24 (2009)
- Drachsler, H., Hummel, H., van den Berg, B., Eshuis, J., Waterink, W., Nadolski, R.J., Berlanga, A.J., Boers, N., Koper, E.J.R.: Effects of the ISIS Recommender System for navigation support in self-organised Learning Networks. *Journal of Educational Technology and Society*, **12**(3), 122–135 (2009b)

- Goldberg, K., Roeder, T., Gupta, D., Perkins, C., Eigentaste, A.: *Constant Time Collaborative Filtering Algorithm*, Vol. 4, pp. 133–151 (New York: Springer 2001)
- Gross, J.L., Yellen, J.: *Graph Theory and Its Applications* (Boca Raton, FL: Chapman & Hall/CRC 2006)
- Herlocker, J.L., Konstan, J.A., Borchers, A., Riedl, J.: Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems* **22**(1), 5–53 (2004)
- Herlocker, J.L., Konstan, J.A., Riedl, J.: Explaining collaborative filtering recommendations. *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, pp. 241–250, 2000, Philadelphia, PA
- Hylén, J.: Open educational resources: opportunities and challenges. *Proceedings of Open Education*, pp. 49–63, September 27–29, 2006, Utah State University, Logan, UT [Retrieved October 5, 2008 from <http://cosl.usu.edu/conferences/opened2006/docs/opened2006-proceedings.pdf/>]
- Keselman, H.J., Huberty, C.J., Lix, L.M., Olejnik, S., Cribbie, R.A., Donahue, B., Kowalchuk, R.K., Lowman, L.L., Petoskey, M.D., Keselman, J.C.: Statistical practices of educational researchers: an analysis of their ANOVA, MANOVA, and ANCOVA analyses. *Review of Educational Research* **68**(3), 350–386 (1998)
- Linden, G., Smith, B., York, J.: Amazon.com recommendations: item-to-item collaborative filtering. *IEEE Internet Computing* **7**(1), 76–80 (2003)
- Longworth, N.: *Lifelong Learning in Action – Transforming Education in the 21st Century* (London: Kogan Page 2003)
- Melville, P., Mooney, R.J., Nagarajan, R.: Content-boosted collaborative filtering for improved recommendations. *Proceedings of 18th National Conference on Artificial Intelligence*, pp. 187–192, July 28–August 1, 2002, Edmonton, Alberta, Canada
- Nadolski, R., Van den Berg, B., Berlanga, A., Drachsler, H., Hummel, H., Koper, R., Sloep, P.: Simulating light-weight personalised recommender systems in learning networks: A case for pedagogy-oriented and rating-based hybrid recommendation strategies. *Journal of Artificial Societies and Social Simulation (JASSS)*, **12**(1), 4 (2009)
- O’Sullivan, D., Wilson, D., Smyth, B.: Improving case-based recommendation a collaborative filtering approach. *Proceedings of 6th European Conference on Case-Based Reasoning, ECCBR*, Aberdeen, Scotland, UK: LNAI Springer, 2002
- Pazzani, M.J.: A framework for collaborative, content-based and demographic filtering. *Artificial Intelligence Review* **13**(5), 393–408 (1999)
- Sarwar, B.M., Karypis, G., Konstan, J., Riedl, J.: Recommender systems for large-scale e-commerce: scalable neighborhood formation using clustering. *Proceedings of Fifth International Conference on Computer and Information Technology*, 2002 [Retrieved October 5, 2008 from internal-pdf://2002_RS_for_large_scale_e-commerce_sarwar_cluster-1331033344/2002_RS_for_large_scale_e-commerce_sarwar_cluster.pdf]
- Schafer, J.B., Konstan, J., Riedl, J.: Recommender systems in e-commerce. Paper presented at the *1st ACM Conference on Electronic Commerce*, Denver, Colorado, 1999
- Suell, T.J.: Designing instructional computing systems for meaningful learning. In M. Jones, P.H. Winne (Eds.): *Adaptive Learning Environments*, pp. 19–54 (Berlin: Springer-Verlag 1992)
- Soboro, I.M., Nicholas, C.K.: Combining content and collaboration in text filtering. *Proceedings of the IJCAI Workshop on Machine Learning in Information Filtering*, pp. 86–91, August 1999, Stockholm, Sweden, 2002
- Tabachnick, B.G., Fidell, L.S.: *Using Multivariate Statistics*, 4th ed. (Boston, MA: Allyn and Bacon 2001)
- Thorpe, M.: *Evaluating Open and Distance Learning* (Essex, UK: Longman Harlow 1988)
- Tisue, S., Wilensky, U.: NetLogo: a simple environment for modeling complexity. *Proceedings of International Conference on Complex Systems*, Boston, MA, USA, 2004
- Vygotsky, L.S.: *Mind in Society: The Development of Higher Psychological Processes* (Cambridge: Harvard University Press 1978)
- Wasserman, S., Faust, K.: *Social Network Analysis: Methods and Applications* (Cambridge: Cambridge University Press 1999)

An overview of earlier work on the application of recommender systems in formal e-learning is given in Nadolski et al. (2008). PRS are already successfully employed in some formal e-learning domains (Andronico et al. 2003; Farzan and Brusilovsky 2006; Tang and McCalla 2004a, b). However, their usefulness for LN is questionable as most of existing PRS heavily rely on rather specific and intensive data provisioning, data maintenance and data-mining, making them ‘heavy-weight’ systems. On the contrary, PRS in LN ask for a ‘light-weight’ and more generally applicable approach. Such a light-weight approach should minimise effort on behalf of the participants, and take into account that LN do not have clear boundaries or structures like in more formal learning settings. Therefore it is not always possible to identify best next LAs for each learner at any time beforehand, as newly added LAs change the LN and enable other learning paths which were formerly unknown. In these, constantly changing LN, it would be too time consuming and/or practically infeasible to provide all necessary – standardised – information beforehand in order to cater for personalised recommendations. For the same reasons, recommendations provided by human intervention are not an option either.

Consequently, PRS for LN should be designed differently from already existing PRS for formal e-learning. To provide practically feasible useful personalised recommendations in LN it is proposed to use (a) a limited registration of the behaviour of the current learner and their peers including rating (i.e., collective behaviour) in the LN and (b) a limited set of LA-characteristics and learners’ characteristics. In other words, it is proposed to use light-weight hybrid PRS (which combines content-based and collaborative filtering). Mainly bottom-up recommendation strategies (RS) seem to be feasible in professional development with a changing and large number of LAs. This is because such collaborative filtering strategies require nearly no maintenance and improve through the emergent behaviour of the community (Hummel et al. 2007).

We intend to provide individualised support for each professional in a LN to increase their goal attainment and satisfaction, and to minimise their study time. We call such support ‘sound recommendations’. In order to arrive at designing them, the key question of our research is:

What RS and which limited set of LA-characteristics and learners’ characteristics is needed in a light-weight hybrid PRS to enable sound recommendations within LNs, and which learner behaviour minimally needs to be traced?

Sound recommendations enable more learners to achieve their goal (i.e., graduating) in less time and with more satisfaction. In other words, we strive for more, faster, and more attractive graduation.

This chapter focuses on a simulation study – consisting of three steps – that addressed this question and that intends to provide insight into how specific RS and PRS key variables affect learning outcomes. It took real data and findings of preceding studies as a starting point (Berlanga et al. 2007; Drachler et al. 2008b). We present three consecutive incremental conceptual simulation models that are based upon a well established approach towards social science simulations (Gilbert and Troitzsch 1999) and are largely in line with Koper’s model of Learning Networks (Koper 2005), their implementation, the enabled recommendation strategies, and

the results obtained from various simulation runs. In this introduction we briefly motivate our approach towards PRS in LN.

8.1.1 PRS in e-Learning

Chapters 6 and 7 argued that application of PRS in e-learning is different from other domains (McCalla 2004). It is beneficial for PRS in LN if pedagogical rules derived from educational psychology research are applied (Koper and Olivier 2004). Since PRS in e-learning are meant to support the learning process, the RS should consist of relevant *pedagogical rules* describing pedagogy-oriented relations between learners' characteristics and LA-characteristics. As an example, we know that from Vygotsky's 'zone of proximal development' (Vygotsky 1978) follows the pedagogical rule 'recommended learning actions with a level just over learners' current competence level'. Other pedagogical rules are: 'go from more simple to more complex', 'learners' effort will increase if they get more satisfied'. Pedagogical rules imply the availability of specific metadata for LAs and of the up to date registration of (a minimal set of) learners' characteristics. Although not specific for PRS in e-learning, also the characteristics of the LN itself (e.g., number of learners, number of LAs, number of sub domains) should be considered as these can effect the impact of PRS on learning outcomes.

Ideally, PRS in e-learning assist learners in finding LAs matching their profile (competences and preferences), keeping them motivated and enabling them to complete their LAs in an effective and efficient way. Purely model-based (i.e., ontology-based) recommendations are very costly and practically infeasible as the required standardised identification at LA-changes, and learner profile changes in the LN cannot be automatised. Although Collaborative Filtering (CF) can address those two problems, purely CF-based PRS systems have severe shortcomings (Van Setten 2005), especially the cold-start problem (new users, new items, and scarcity of past user actions) which can be overcome by using hybrid PRS system, that also use (a limited number of) LA-characteristics and learner-characteristics. Hybrid PRS systems have been shown to outperform purely CF-based PRS systems (Claypool et al. 1999; Good et al. 1999; Melville et al. 2002; Pazzani 1999; Soboro and Nicholas 2002). Tang and McCalla (2004b) performed a simulation study claiming that hybrid PRS systems can even be as effective as purely model-based PRS.

Although purely model-based PRS mostly outperform other PRS-systems, they are not appropriate for LN with a fast changing and potentially huge number of LAs. Furthermore, a hybrid PRS with intensive model-based data maintenance and data-mining CF techniques would also be impractical as they induce enormous network traffic (distributed data) and require huge computing power. What instead is needed is a light-weight hybrid PRS with minimised data provisioning, data maintenance, and data-mining.

8.1.2 Simulations for PRS in e-Learning

Although the field experiment (see Chap. 7) showed promising results, at the same time various practical constraints (e.g., limited number of learners and LAs) made it difficult to investigate other treatments such as other hybrid RS with or without rating.

If we want to develop a PRS for learners in a LN we face the problem that limited real data are currently available from user studies. On the other hand, general findings with respect to pedagogical rules, previous applications of PRS in e-learning, our own field experiment, and expert reasoning can assist us in articulating a conceptual simulation model for PRS. Simulations can support defining requirements of a PRS for LN before actually starting the costly process of development, implementation, testing and revision in real education. Field experiments with real learners need careful preparation as they cannot be easily repeated or adjusted within a condensed timeframe. Another advantage of simulations is that they bypass some ethical and practical constraints of field experiments.

Although simulations are a well established approach in social science (Gilbert and Troitzsch 1999), they have been sparsely used for recommendations in e-learning and LN. Tang and McCalla (2004b) performed a simulation study showing no differences between model-based and hybrid-based paper recommendations including rating. Koper (2005) described a simulation study showing that selecting units of learning (i.e., a structured set of learning actions) informed by indirect social interaction, increased learner retention in a LN when compared to a selection without indirect social interaction. Inclusion of CF-based selection of units-of-learning appeared beneficial for learner retention. Koper asked for solutions to improve the RS (i.e., decrease the matching error). As a possible solution to decrease the matching error we suggest using hybrid techniques with ratings to improve recommendations as they enable a light-weight PRS, also based on findings in an earlier simulation study (Berlanga et al. 2007). To the best of our knowledge, no study specifically addresses the effect of rating towards recommendations in LN. This is rather surprising as many previously described hybrid PRS included rating (Tang and McCalla 2004a, b; RACOFI, CourseAgent).

8.2 Conceptual Simulation Models and Setting Up Simulations

As the study consists of three incremental steps – each using a conceptual simulation model and a Netlogo (Wilensky 1999) implementation – three separate method and results sections will be distinguished. A method section describes hypotheses for the simulation study and its setup of each step. Netlogo is used to perform the simulation runs. The outcomes of such runs validate the version of the conceptual simulation model in that step, and guide further articulation of the conceptual simulation model. Our study started with a focussed expert discussion in which variables, their relationships, and associated literature with existing research were identified as input for all steps. In the first step, focus is given towards simplicity. The second step

focuses on more reality, whereas the third step focuses on practical implementation issues. Each model will be elaborated in three parts: model variables, measurement variables, and recommendation strategies, followed by the setup of the simulation. Finally, conditions and treatments that are included in the simulation runs of that step will be described. Each conceptual simulation model is largely in line with the model described by Koper (2005). The main goal of the final version of the conceptual simulation model is that it represents the minimised set of LA- and learner characteristics that can support the included RS.

After such a method section, the result section will recapitulate the main results of the simulation runs for that step. A step-specific discussion will be restricted to implications for the following step. The general discussion at the end of this chapter will focus on the results from the last step, and identifies preferred RS in view of our key question, describes some limitations of this study, and provides some suggestions for future research.

8.2.1 Method Step 1: Benefits of Hybrid RS

During this step, the main aim was to verify whether hybrid recommendation strategies in e-learning were in favour to purely CF-based RS. Please note that purely model-based RS (i.e., ontology-based) do not comply with our main goal of light-weight PRS although it might be expected that ontology-based recommendations will address the aim of more, faster, and more attractive graduation most closely. Graduation is defined as reaching the learning target i.e. completing a set of LA's successfully. Our two main hypotheses for this step are:

H1: PRS recommendations yield more, more satisfying, and faster graduation than no recommendations

H2: hybrid recommendations from PRS yield more, more satisfying, and faster graduation than purely CF-based RS.

The conceptual simulation model describes variables, their initial value distribution, as well as their relationships, which are often represented by formulas. This model is used as input for implementation within a simulation environment, i.e., Netlogo 4.02 and takes into account learner profiles for the current learner and his or her peers as well as the LA-characteristics. The RS produces a personalised recommendation of a best next LA. Individual learner behaviour is modelled in the learner model. It is indirectly influenced by peers if CF is included in the RS during setup, which is dealt with in the run model. The amount of alignment between learner profile and LA-characteristics, as well as the current state of the learners model (for example: learner competence) influence the chance of successful LA-completion. The RS supports better alignment between learner profiles and LA-characteristics, whereas the learner model represents learners' changing behaviour when trying to achieve goals.

Figure 8.1 shows an integrated picture of the program flow in the simulation environment (the thick lines and arrows) and the conceptual simulation model. Details

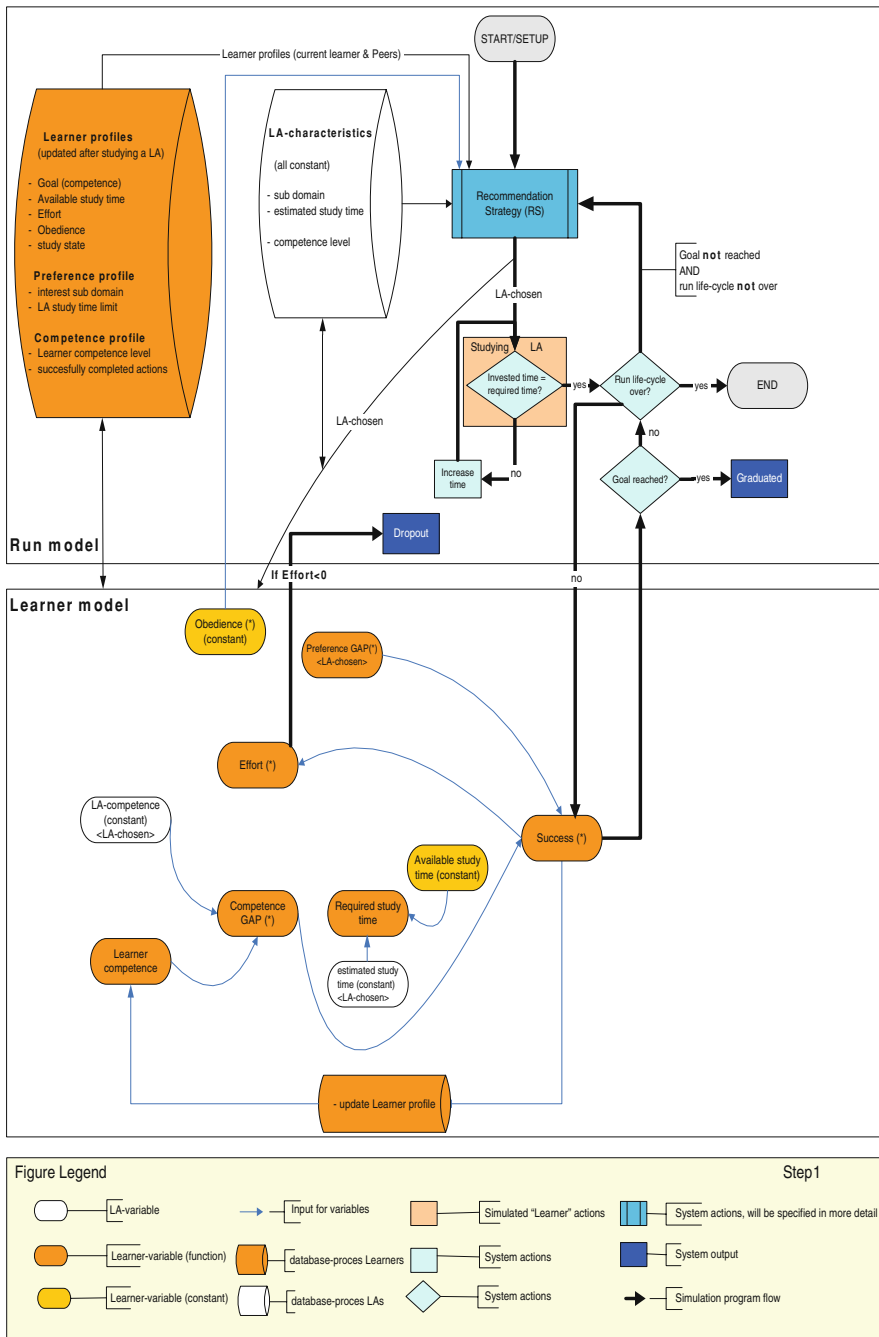


Fig. 8.1 Conceptual simulation model and simulation program flow (Step 1)

about the RS are included in Fig. 8.2. Details about the setup are presented in Fig. 8.3.

Table 8.1 presents an overview of all variables in the conceptual simulation model in Step 1 and their implementation within the simulation. Some variables are related by formulas and are further detailed out in Table 8.2 describing pedagogical rules in the conceptual simulation model.

Three measurements were used to test this study's hypotheses.

Graduation. Learners will dropout if their effort falls below zero and consequently fail to reach their goal. Reaching their goal equals graduating. Identifying effectiveness of the RS from PRS will be based upon the percentage of learners reaching their goal (Graduates). A higher percentage of graduates indicates more effectiveness.

Satisfaction. Satisfaction is measured when learners achieve their goal. We suppose that a higher proportion with maximum satisfaction (or more effort) at goal completion indicates that they are more satisfied at graduation than a lower proportion with maximum satisfaction (or less effort) at graduation. Learners' willingness to voluntarily invest a certain effort to study is here regarded to be similar to satisfaction.

Time to graduate. Learners reaching their goal graduate. The impact of RS from PRS on time efficiency is determined by identifying learners' total study time for achieving their goal: time to graduate. The less time to graduate, the more time efficient.

Application of the RS ultimately results in a LA-chosen by the learner (Fig. 8.2). The simulation distinguishes three treatment groups each using a different RS and one control group. The RS can partly be defined in the simulation setup. Recommendations can be based upon: Ontology (O), Peers (P), Ontology and Peers (OP), resulting in three treatment groups.

Figure 8.1 also shows the simulation program flow. During the setup the following characteristics can be considered/defined (see Fig. 8.3):

- the Recommendation Strategy (RS)
- the Learning Network characteristics (number of learners, number of LAs, number of sub domains which also determines number of interests at the learner side, and the number of LAs for each competence level in the target Goal)
- other, general settings of the simulation (e.g., run length)

Considering these values, the setup initialises the environment as follows:

- the number of sub domains is randomly uniformly distributed across LAs
- each LA is initialised with a random estimated study time, sub domain
- each learner is initialised with an uniformly distributed random: interest sub domain, LA study time limit (preference b), available study time
- each learner starts with learner competence and will have the same goal
- each learner starts with a normally distributed random effort ($M=10$, $SD=10/3$)
- each learner has a specific obedience, normally distributed ($M=0.6$, $SD=0.15$)

[Explanation Fig. 8.1, see also Table 8.1 and 8.2 and Fig. 8.2]

Each [Learner profile] consists of a [Preference profile], a [Competence profile] and some other learner characteristics (like goal or available study time), all being used within the RS. Preferences a and b in [Preference profile] are 'interest sub domain' and 'LA study time limit'.

[Competence profile] is restricted to one competence which can include three levels for the [Goal]. Successfully completed LAs contribute to this level. In the simulation setup, the number of successfully completed LAs for this level can be specified. For simplification of the simulation, learners start with the same [Learner competence level] and have the same goal. [Learner competence level] indicates the learner's achievement with respect to the goal and is updated by the outcome of [Success].

[Available study time] has the same magnitude as the simulation tick time (1 tick = 1 week).

[Effort] is initially normal distributed amongst learners, but changes during study. Effort is the key variable which determines dropout. If Effort gets below zero, a learner will drop out the LN and will not graduate. Effort depends on previous effort and [Success].

[Preference GAP] measures alignment between learners' preferences and corresponding LA-characteristics. Here, learners' preferences are restricted to [interest sub domain] and [LA study time limit]: the maximum time a learner wants to invest for studying an LA and before doing its associated LA-examination.

[Competence GAP] measures alignment between [Learner competence level] and [LA-competence level] of the LA-chosen. Please note that [Learner competence level] is a variable with a specific value at a certain point of time, whereas [LA-competence level] is a constant. Mismatches for preferences and/or competences will decrease the chance for Success, whereas better matches or preferably perfect matches will increase the chance for Success.

[Obedience] is normally distributed amongst learners but remains constant for each learner and refers to users' willingness to obey the recommendation.

[Required study time] is used in the simulation variable [required time]: the quotient of required study time and simulation tick time.

If [Success] is true, the learner passes LA-examination and [learner competence] will be adjusted, whereas [effort] will increase. Note that effort in itself does not influence the chance of success in this version of the model. The match in [Success] considers if the competence level and the estimated study

load of the advised learning activity are, correspondingly, smaller or equal to the learner's competence level and to the learner's LA study time limit, and if the sub domain of the LA is the same as the learner's sub domain interest. The LA-examination fails (i.e., [Success] is false) if there is no perfect match on all three criteria.

All LA-characteristics remain unchanged in this version of the model.

Every condition in the simulation was replicated 12 times (i.e., $N = 12$ runs) to justify the use of classical statistic techniques on resulting data (Law and Kelton 2000, p. 496), analyzed with SPSS version 15. The data set of the simulation outcomes can be found at <http://hdl.handle.net/1820/1491>. Each condition included three treatment groups (O, P, OP) and one control group (C). Six conditions were included: goal (2) \times sub domain (3). A low-end goal and a high-end goal were taken (Tables 8.3 and 8.4). Either, 2, 3 or 4 sub domains were taken as they were all regarded to be representative for studies in higher distance education. In all six conditions, each group included 250 learners that could choose from 400 LAs. For all runs, only 'graduates' or 'dropouts' were allowed after run length. In other words, no participants were 'still studying' after run length. To check the robustness of the conceptual simulation model other conditions have been simulated as well. However, those will only be sparsely reported upon as these conditions showed similar results to the ones included.

Results Step 1

Hypothesis 1 (PRS recommendations yield more, more satisfied, and faster graduation than no recommendations) was tested by using regular statistical methods, such as analyses of variance on a global level, continuing with Bonferroni's correction when using multiple comparisons at a more detailed level. Here, multiple comparisons were always conducted between the control group and one of the three treatment groups. Data means and standard deviations for all three measurements for all conditions are presented in Tables 8.3 and 8.4.

Graduation. Analyses of variance showed – for all six conditions a significant difference in the percentage of Graduates between all groups. More detailed analysis (Table 8.5) showed that the control group always had a significant smaller percentage of graduates than all other treatment groups. This confirms Part 1 of H1 that PRS recommendations yield to more graduation than no recommendations.

Satisfaction-effort. Analyses of variance showed for all six conditions a significant difference in effort at graduation between all groups. More detailed analysis (Table 8.6), namely multiple comparisons all with Bonferroni's cor-

Table 8.1 Overview variables in conceptual simulation model and their implementation within the simulation (Step 1). See Table 8.2 for the formulas

Variable	Description	Implementation in simulation		
		Range (initialisation)	Formula	Input for
Learner profile	–	–	–	RS
– Goal	Competence level(s)	In set up, same for all, allows needed number of LAs	No	
– Available study time	–	M = 20 hours/week, SD = 5 (Normally-distributed)	No	Required study time
– Effort	Investment to study	M = 10, SD = 10/3 (Normally-distributed)	Yes	Study state
– Obedience	Follow up	M = 0.6, SD = 0.15 (Normally-distributed)	No	RS
– Study state	recommendation	[studying (in progress), graduated, dropout]		Dropout, graduated
– Preference profile		–	–	
– Interest sub domain (a)		Number of sub domains, in setup (Randomised)	No	Preference GAP
– LA study time limit (STL) (b)	–	[5, 10, 15] (Randomised)	No	Preference GAP
	–	–	–	Competence GAP
– Competence profile		0, updated if number of successfully completed LAs	Yes	Competence level
– Learner competence level (LCL)	–	matches level	Yes	
	Contribution towards goal	An integer for each applicable level in the goal		
– Successfully completed LAs				
LA characteristics	–	–	–	RS
– Sub domain (flavour a)	–	Number of sub domains in setup (Randomised)	No	Preference GAP
– Estimated study time (flavour b)	–	[50, 100, 150] hours (Randomised)	No	Required study time
– Competence level (CL)			No	
– Preference GAP	Alignment preferences and flavours	Integer [–2, 0]	Yes	Success
– Competence GAP	Alignment LCL and CL of LA	Integer [–1, 0]	Yes	Success
– Required study time (rst)	Invested time before LA-examination	Same as estimated study time (constant and independent!)	No	Studying LA
– Success	Learner passes or fails LA-examination	Boolean	Yes	Effort, LCL

Table 8.2 Formulas and description for all variables with changing values in conceptual simulation model (Step 1)

Variable	Description	Formula in simulation	Input for
Learner profile – Competence profile – Learner competence level – Successfully completed LAs	Goal accomplishment Contribution towards goal	A specific LA competence level is mastered if the number of successfully completed LAs with a specific LA competence level matches the corresponding definition in the goal (specified at the setup). The number of successfully completed LAs for a specific LA competence level is stored within ‘successfully completed LAs’.	Competence GAP
– Effort	Investment to study	$effort = previous\ effort + 0.5$ (if Success = 0, true), $effort = previous\ effort - 1.0$ (if Study state Success < 0, false)	
– Preference GAP	Alignment preferences and flavours	Preference GAP = – SUM (Learner preference(x_i) – LA preference(x_i)), $i = 1, 2$ Learner preference(x_i) – LA preference(x_i) = 1 if Learner preference(x_i) \neq LA preference(x_i) Learner preference(x_i) – LA preference(x_i) = 0 if Learner preference(x_i) = LA preference(x_i) The SUM is 0 if the preference GAP is 0, indicating that there is a perfect match. For calculation purposes (in Success) we use negative values if there is no perfect match.	Success
– Competence GAP	Alignment learner competence level and competence level of LA	Competence GAP = (Learner competence level – LA competence level) + 1 learner competence level: 0, 1, 2 LA competence level: 1, 2, 3 For calculation purposes a symmetric distribution of competence GAP with integers ranging from [–2, 2] was preferred. Therefore, a perfect match results in a competence GAP of 0.	Success
– Success	Learner passes or fails LA-examination	Success = preference GAP + competence GAP The learner successfully completes a LA if Success = 0, otherwise the LA examination fails.	Effort, LCL

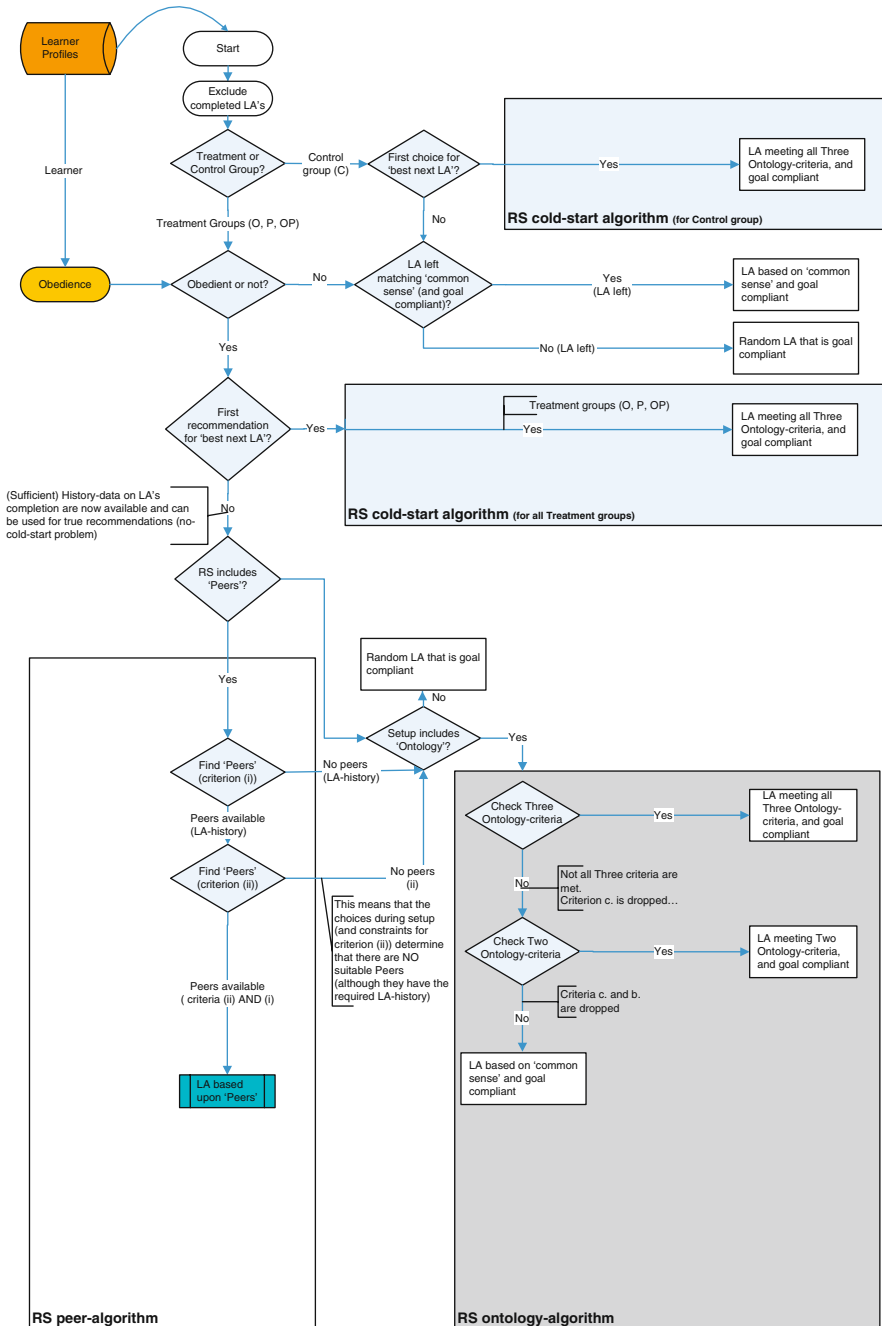


Fig. 8.2 Recommendation strategy (RS) for the simulation program (Step 1 and 2)

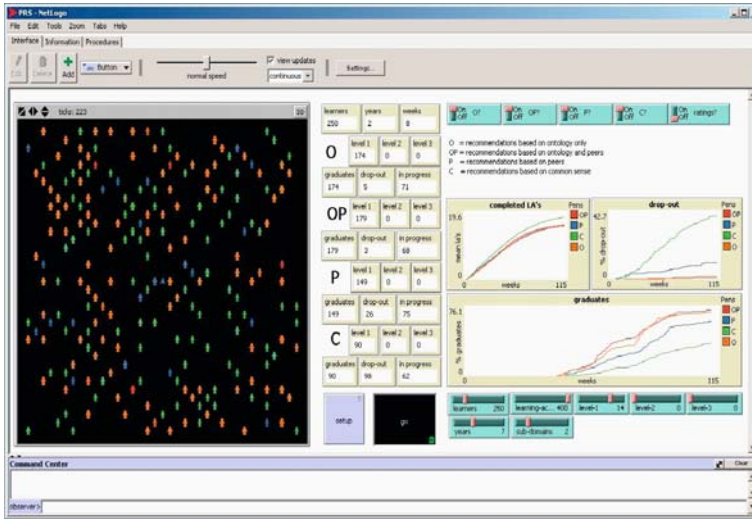


Fig. 8.3 Screenshot of the simulation in Netlogo. Level-1 goal: 14 LAs level-1, 2 sub-domains 400 LAs, 250 learners for each group, run length 7 years

Table 8.3 Graduation, satisfaction, time to graduate for a low-end goal (Step 1)

# Sub domains	Variables	Treatment						No treatment Control (C)	
		Ontology (O)		Peers (P)		OP		M	SD
2	Graduates	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
	– perc.	87.5	2.5	83.6	3.1	87.8	2.1	27.2	3.1
	– satisf.	6.9	3.5	6.3	3.4	6.9	3.4	3.3	2.8
3	Graduates								
	– perc.	81.5	5.5	77.7	7.0	81.9	5.2	27.2	3.1
	– satisf.	6.7	3.5	6.1	3.3	6.7	3.5	3.8	2.6
4	Graduates								
	– perc.	67.5	4.8	65.0	7.5	67.6	5.6	12.4	2.7
	– satisf.	6.6	3.4	5.9	3.3	6.6	3.5	3.8	2.6
	– time	1478	359	1555	357	1470	355	2154	355

Notes: number of learners = 250 (for each group), number of LAs = 400, run length = 7 year, 12 runs.

Goal: 14 LAs level 1. Perc. = percentage graduates. Satisf. = effort at graduation (for Step 1) or percentage graduates with maximum satisfaction (Step 2 and 3). Time = time to graduate.

reaction, showed that the control group always had lower effort at graduation than any treatment group. This confirms H1 (Part 2) that PRS recommendations yield to more satisfied graduation than no recommendations.

Table 8.4 Graduation, satisfaction, time to graduate for a high-end goal (Step 1)

# Sub domains	Variables	Treatment						No treatment Control (C)	
		Ontology (O)		Peers (P)		OP		<i>M</i>	<i>SD</i>
2	Graduates	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
	– perc.	32.6	4.1	22.2	2.3	30.8	5.4	0.2	0.3
	– satisf.	5.3	3.3	4.1	2.8	5.0	3.2	1.7	1.1
3	Graduates								
	– perc.	5755	969	5994	877	5687	930	7472	319
	– satisf.								
4	Graduates								
	– perc.	20.8	4.3	11.6	2.5	19.3	3.5	0.0	0.0
	– satisf.	4.9	3.2	3.6	2.4	4.7	3.1	–	–
4	Graduates								
	– perc.	6194	883	6624	638	6207	846	–	–
	– satisf.								
4	Graduates								
	– perc.	14.2	2.6	7.1	2.2	12.9	2.4	0.0	0.0
	– satisf.	4.6	3.0	3.5	2.3	4.8	3.0	–	–
	– time	6385	612	6690	501	6384	584	–	–

Notes: number of learners = 250 (for each group), number of LAs = 400, run length = 15 year, 12 runs.

Goal: 14 LAs level 1, 28 LAs level 2, 14 LAs level 3. Perc. = percentage graduates. Satisf. = effort at graduation (for Step 1) or percentage graduates with maximum satisfaction (Step 2 and 3). Time = time to graduate.

Table 8.5 Outcomes for (a) Analyses of variance and (b) Multiple comparisons with Bonferroni's correction with respect to graduation (N = 12 runs) (Step 1)

Goal	# Sub domains	Analyses of variance			Multiple comparisons
		<i>F</i>	<i>MSE</i>	<i>p</i>	
Level-1	2	$F(3, 11996) = 1905.00$	1052	<0.05*	C fewer <i>G</i> than any T
	3	$F(3, 11996) = 1992.78$	1219	<0.05*	C fewer <i>G</i> than any T
	4	$F(3, 11996) = 1143.79$	886	<0.05*	C fewer <i>G</i> than any T
Level-3	2	$F(3, 11996) = 435.74$	265	<0.05*	C fewer <i>G</i> than any T
	3	$F(3, 11996) = 255.99$	108	<0.05*	C fewer <i>G</i> than any T
	4	$F(3, 11996) = 167.56$	50	<0.05*	C fewer <i>G</i> than any T

Notes: Multiple comparisons (mean difference between two groups, all $p < 0.05^*$). C = control group, T = treatment group, *G* = Graduates.

Time. Analyses of variance showed for all six conditions a significant difference in time to graduate between all groups. More detailed analysis, namely multiple comparisons all with Bonferroni's correction, showed that the control group always needed significantly more time to graduate than any treatment group (Table 8.7). This confirms H1 (Part 3) that PRS recommendations yield to faster graduation than no recommendations.

Hypothesis 2 (hybrid recommendations from PRS yield more, more satisfied, and faster graduation than purely CF-based RS) was also tested by using regular statistical methods, namely Bonferroni's correction when using multiple comparisons at a more detailed level, using mean difference between two groups (P-group

Table 8.6 Outcomes for (a) Analyses of variance and (b) Multiple comparisons with Bonferroni's correction with respect to satisfaction at graduation (N = 12 runs) (Step 1)

Goal	# Sub domains	Analyses of variance <i>F</i>	<i>MSE</i>	<i>p</i>	Multiple comparisons
Level-1	2	$F(3, 8578) = 146.52$	1675	<0.05*	C fewer <i>S</i> than any T
	3	$F(3, 7731) = 113.49$	1312	<0.05*	C fewer <i>S</i> than any T
	4	$F(3, 6367) = 85.74$	970	<0.05*	C fewer <i>S</i> than any T
Level-3	2	$F(3, 2572) = 22.35$	214	<0.05*	C fewer <i>S</i> than any T
	3	$F(3, 1548) = 22.02$	198	<0.05*	C fewer <i>S</i> than any T
	4	$F(3, 1022) = 16.30$	133	<0.05*	C fewer <i>S</i> than any T

Notes: Multiple comparisons (mean difference between two groups, all $p < 0.05^*$). C= control group, T = treatment group, *S* = satisfaction at graduation. Different 'n'-s in the F-statistics as conditions differ in number of graduates.

Table 8.7 Outcomes for (a) Analyses of variance and (b) Multiple comparisons with Bonferroni's correction with respect to time to graduate (N = 12 runs) (Step 1)

Goal	# Sub domains	Analyses of variance <i>F</i>	<i>MSE</i>	<i>p</i>	Multiple comparisons
Level-1	2	$F(3, 8578) = 695.82$	95850732	<0.05*	C more <i>T</i> than any T
	3	$F(3, 7731) = 572.38$	81655702	<0.05*	C more <i>T</i> than any T
	4	$F(3, 6367) = 414.17$	52679178	<0.05*	C more <i>T</i> than any T
Level-3	2	$F(3, 2572) = 22.49$	19486840	<0.05*	C more <i>T</i> than any T
	3	$F(3, 1548) = 36.19$	24312766	<0.05*	C more <i>T</i> than any T
	4	$F(3, 1022) = 23.39$	7863660	<0.05*	C more <i>T</i> than any T

Notes: Multiple comparisons (mean difference between two groups, all $p < 0.05^*$). C= control group, T = treatment group, *T* = time to graduate. Different 'n'-s in the F-statistics as conditions differ in number of graduates.

and OP-group), $p < 0.05$. These multiple comparisons were always conducted after analyses of variance on the more global level. As before, results for all measurement variables are presented in Tables 8.3 and 8.4.

Graduation. Multiple comparisons, using Bonferroni's correction, revealed – for all six conditions – that the OP-group had significantly more graduates than the P-group, using mean difference between two groups, $p < 0.05$. This confirms Part 1 of H2. Please recall that there was a significant difference in the percentage of graduation between all groups for all six conditions when analysis of variance was used.

Satisfaction – effort. Multiple comparisons, using Bonferroni's correction, revealed – for all six conditions – that the OP-group had significantly more effort at graduation than the P-group, using mean difference between two groups, $p < 0.05$. This confirms Part 2 of H2. Please recall that there was a significant difference in effort at graduation between all groups for all six conditions when analysis of variance was used.

Time. Multiple comparisons, using Bonferroni's correction, revealed – for all six conditions – that the OP-group needed significantly less graduation time

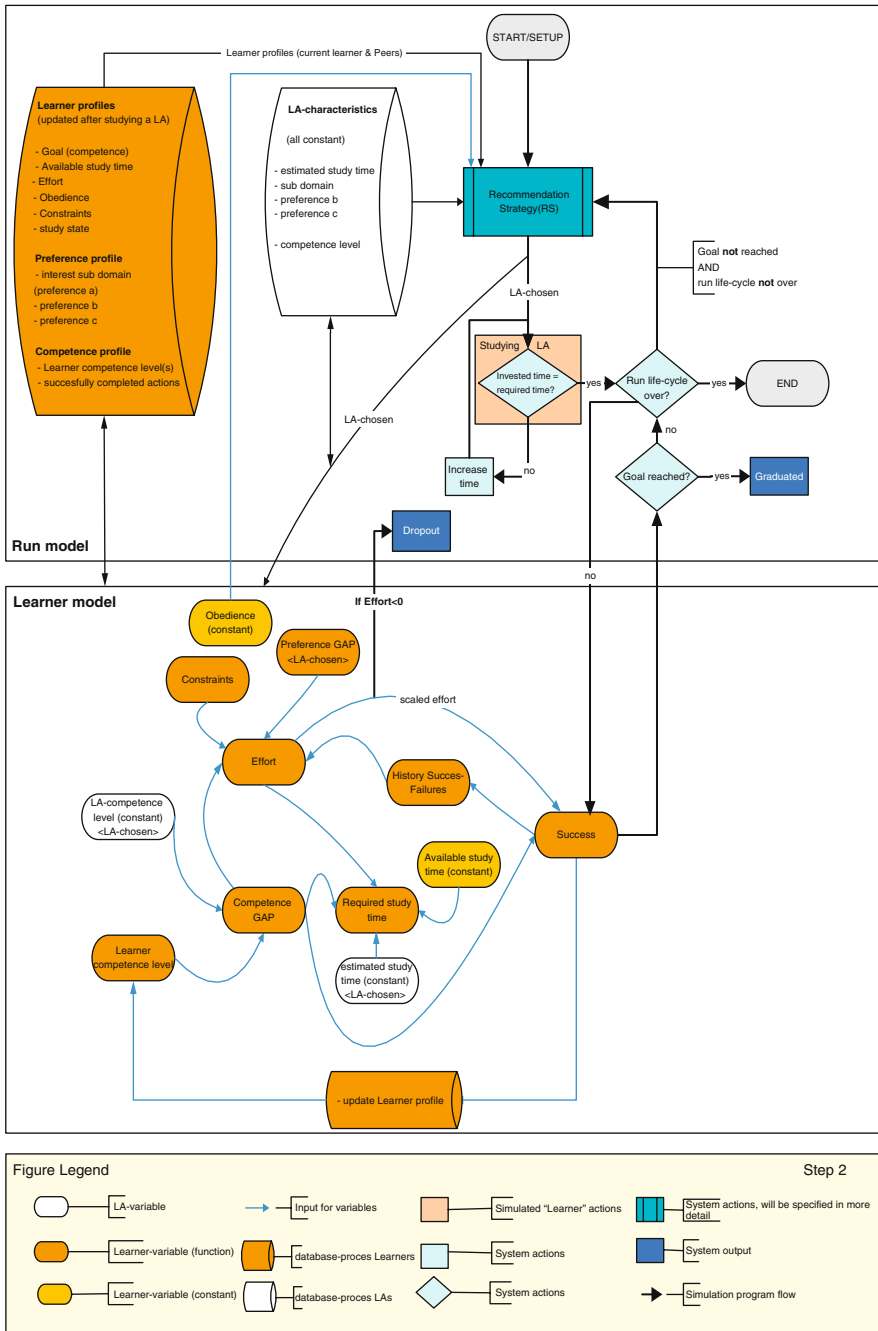


Fig. 8.4 Conceptual simulation model and simulation program flow (Step 2)

than the P-group, using mean difference between two groups, $p < 0.05$. All results were also present at the level of separate runs. This confirms Part 3 of H2. Please recall that there was a significant difference in graduation time between all groups for all six conditions when analysis of variance was used.

The findings in this step demonstrate the advantage of using PRS recommendations in e-learning as these yield to more graduation, faster graduation, and more satisfied graduation. Furthermore, the advantage of applying hybrid recommendations was confirmed. However, in particular for the high-end goal, the concrete figures for graduation seem unrealistic for all groups.

The conceptual simulation model in Step 1 is a strong simplification of the real world. Although pedagogical rules are used, they are scarce and they need more fine-tuning to gain higher realism. In Step 1, [Success] in the learner model was determined by several variables but was represented quite abstract and rather unrealistic. [Success] effected [effort] and as a consequence drop out, but [effort] did not directly influence [Success], which seems more realistic. Furthermore, [effort] is expected to be influenced by other variables as [Success] only, for example by [Competence GAP], [Preference GAP] and [Constraints]. Therefore, for Step 2 it was decided to consider [effort] as – a more concrete – key variable determining [Success]. In this way, a loop between [Success] and [effort] was introduced into the learner model. Both variables were expected to mutually influence each other, which is regarded more realistic. In addition, it was decided that [required study time] should depend on learner variables and not only on the LA-characteristic ‘estimated study time’. Finally, it was decided to include some additional fine-tuning, for example with respect to [Preference profile]. As mentioned before, the focus for Step 2 is realism of the conceptual simulation model.

8.2.2 Method Step 2: Fine-Tuning Hybrid RS

During this step, the main aim was to have a more realistic conceptual simulation model, and again – as with Step 1 – to verify whether hybrid recommendation strategies in e-learning were in favour to purely CF-based RS. Our two main hypotheses for this step are exactly the same as in the previous step.

Figure 8.4 shows an integrated picture of the program flow in the simulation environment (the thick lines and arrows) and the conceptual simulation model for Step 2. The amount of alignment between learner profile and LA-characteristics and the current state of the learners model (for example: learners’ effort) influence the chance of successful LA-completion. The RS supports better alignment between learner profiles and LA-characteristics, whereas the learner model represents learners’ changing behaviour when trying to achieve goals.

Table 8.8 presents an overview of all variables in the conceptual simulation model and their implementation within the simulation. Table 8.9 summarises relations and formulas, describing pedagogical rules in the conceptual simulation model. These rules are used within the RS.

[Explanation Fig. 8.4, see also Tables 8.8 and 8.9 and Fig. 8.2]

Each [Learner profile] consists of a [Preference profile], a [Competence profile] and some other learner characteristics (like goal or available study time), all being used within the RS. Preferences b and c in the [Preference profile] could for example be learning strategy, presentation style, or price to enrol.

[Competence profile] is exactly the same as for Step 1. Again, [Learner competence level] indicates the learner's achievement with respect to the goal and is updated by the outcome of [Success].

[Available study time] has the same magnitude as the simulation tick time (1 tick = 1 week) but has a different distribution amongst learners as before.

[Effort] is initially normally distributed amongst learners, but changes during study (see Table 8.9). Effort is the key variable which determines dropout (Ryan and Deci 2000). If Effort gets below zero, a learner will drop out the LN and will not graduate. Effort depends on previous effort, [Preference GAP], [Competence GAP], [Constraints] and [History of Success/Failures] (abbreviated as: History) on the last three LA examinations.

[Preference GAP] measures alignment between learners' preferences and corresponding LA-characteristics. The smaller the gap, the better the alignment.

[Competence GAP] measures alignment between [Learner competence level] and [LA-competence level] of the LA-chosen. A perfect match occurs if [LA-competence level] is one level above [Learner competence level] (Vygotsky 1978). Please note that [Learner competence level] is a variable with a specific value at a certain point of time, whereas [LA-competence level] is a constant. Mismatches for preferences and/or competences will decrease effort, whereas better matches or preferably perfect matches will increase effort.

[Constraints] (like fatigue, being in the flow, a noisy or quiet or study room, stress) can influence the amount of effort learners want to invest when studying. Constraints are randomised at each studied LA. For calculation purposes, we define constraints as '1' in case of positive effects, '-1' in case of negative effects, and '0' in case of a neutral effect. As constraints are considered to be a multidimensional construct, a more fine grained approach could be used. For the sake of simplicity, we have conceived these constraints to be a one-dimensional construct.

[History] also affects [Effort]. Several successes in a row is expected to increase effort (more motivated), whereas successive failure will be detrimental to learner's effort, and ultimately could result in drop out.

[Obedience] differs between learners but remains constant for each learner. Obedience is similar to predictive utility that measures influences of system predictions upon users' willingness whether or not 'consuming' an item (i.e.,

obeying the recommendation) (Konstan et al. 1997; Walker et al. 2004). In an earlier study we identified an obedience level of 60% (Drachsler et al. 2008b) which is similar to other studies (Bolman et al. 2007; Cranen 2007).

[Required study time] is the time a learner invests before doing an LA-examination. In case of a competence GAP, a learner needs more (if knowledge deficiency) or less time (if knowledge surplus) than the [estimated study time] of the LA-chosen. Required study time is used in the simulation variable [required time]: the quotient of required study time and simulation tick time.

If [Success] is true, the learner passes LA-examination and achievement of the learner competence level corresponding with the LA competence level and goal will improve. If the studied LA is too far above the learners' current Learner competence level (in other words, there is a very huge competence GAP), it does not matter how much effort the learner invest, it will always lead to failing this LA-examination. However, for LA's that normally would be somewhat beyond learners' scope of possibilities, more effort can lead to their successful completion.

As for Step 1, all LA-characteristics remain unchanged in this model.

The measurement variables are the same but satisfaction slightly differs as it now relates to the proportion of graduates with maximum satisfaction. The applied RS at this step is exactly the same as for Step 1 (Fig. 8.2).

During the setup the same characteristics can be considered/defined as at Step 1. As with Step 1, every condition in the simulation for Step 2 was replicated 12 times (i.e., $N = 12$ runs). The data set of the simulation outcomes can be found at: <http://hdl.handle.net/1820/1491>. Each condition included three treatment groups (O, P, OP) and one control group (C). The same six conditions as for Step 1 were included: goal (2) \times sub domain (3). In order to save space, we will only include detailed figures about two conditions (low-end goal for 3 sub domains, and a high-end goal for 3 sub domains) and we will not repeat notes to Tables as they are the same as for Step 1.

Results Step 2

Hypothesis 1 (PRS recommendations yield more, more satisfied, and faster graduation than no recommendations) was tested by using regular statistical methods, such as analyses of variance on a global level, continuing with Bonferroni's correction when using multiple comparisons at a more detailed level. Here, multiple comparisons were always conducted between the control group and one of the three treatment groups. Data means and standard deviations for all three measurements for all conditions are presented in Tables 8.10 and 8.11.

Table 8.8 Overview variables in conceptual simulation model and their implementation with the simulation (Step 2). See Table 8.9 for the formulas

Variable	Description	Implementation in simulation		
		Range (initialisation)	Formula	Input for
Learner profile				
- Goal	- Competence level(s)	-	-	RS
- Available study time	-	In set up, same for all, allows 3 levels and needed number of LAs	No	Required study time
- Effort (and scaled effort)	- Investment to study	M = 20 hours/week, SD = 5 (Normally-distributed)	No	
- Obedience	- Follow up recommendation	M = 10, SD = 3 (Normally-distributed)	Yes	
- Constraints	- Fatigue, flow, stress, a.s.o.	M = 0.6, SD = 0.15 (Normally-distributed)	No	Success/dropout, rst
- Study state	-	[-1, 0, 1] (Randomised) for each studied LA	No	RS
- Preference a (sub domain)	-	[studying (in progress), graduated, dropout]	No	effort
- Preferences b, c	-	number of sub domains, in setup (Randomised)	No	Dropout, graduated
- Learner competence level (LCL)	- Goal accomplishment	[b1, b2, b3] (Randomised), [c1, c2, c3] (Randomised)	Yes	Preference GAP
- Successfully completed LAs	- Contribution towards goal	0, updated if number of successfully completed LAs matches level	Yes	Preference GAP
		an integer for each applicable level in the goal		Competence GAP
LA characteristics				
- Estimated study time	-	-	-	Competence level
- Sub domain (flavour a)	-	100 hours for each LA	-	RS
- Preferences b,c (flavours b,c)	-	number of sub domains in setup (Randomised)	No	Required study time
- Competence level (CL)	-	[b1, b2, b3] (Randomised), [c1, c2, c3] (Randomised)	No	Preference GAP
- Preference GAP	-	same as level(s) in goal, in setup (Randomised)	No	Preference GAP
	- Alignment preferences and flavours	Integer [-3, 0]	Yes	Competence GAP
- Competence GAP	- Alignment LCL and CL of LA	Integer [-2, 2]		Effort
	- Invested time before LA-examination	[50, 100, 150]	Yes	Effort, Success, rst
- Required study time (rst)	- Learner passes or fails LA-examination	Boolean	Yes	Studying LA
- Success			Yes	Effort, LCL

Table 8.9 Formulas and description for all variables with changing values in conceptual simulation model (Step 2)

Variable	Description	Formula in simulation	Input for
- Learner competence level	Goal accomplishment Contribution towards goal	Exactly the same as for Step 1	Competence GAP
- Successfully completed LAs	Investment to study and satisfaction during study	Effort = PE + SUM ($w_1 * PG, w_2 * G(CG), w_3 * Constraints, w_4 * History Success/Failures$) - PE = Previous Effort; PG = Preference GAP; CG = Competence GAP - $w_1 = w_2 = w_3 = w_4 = 1$ (all weighting values); PG: -3, -2, -1, 0; CG: -2, -1, 0, 1, 2 - G(CG): G(-2)=-1, G(-1)=1, G(0)=0, G(1)=-1, G(2)=-2, → output [-2, 1]; - Constraints: -1, 0, 1 - History Success/Failures (abbreviated as: History) for each LA: -1, 0, 1 →, output [-3, 3]	Success, dropout, required study time
- Effort		Effort is scaled in order to be able to deal with different weighting values for its input variables (all being 1 for this study) and is scaled for calculating Success: Scaled Effort (SE). IF $0 \leq Effort < 7 \rightarrow Scaled\ Effort = 1$; IF $7 \leq Effort <= 13 \rightarrow Scaled\ Effort = 2$; IF $14 \leq Effort <= 20 \rightarrow Scaled\ Effort = 3$	
- Preference GAP	Similar to Step 1	Similar as for Step 1, however here three preference variables are used	Effort

Table 8.9 (continued)

Variable	Description	Formula in simulation	Input for
- Competence GAP	See Step 1	Exactly the same as for Step 1	Effort, Success, rst
- Required study time (rst)		$\text{rst} = \text{FSA} * ((1 + \text{H}(\text{CG}*)) \text{LA estimated study time})$ - FSA = Factor Scaled Effort: FSA = 0.8 if SE = 3, FSA = 1.0 if SE = 2, FSA = 1.2 if SE = 1 - H(CG) is a function with competence GAP (CG) as input - H(CG) = 0.5 if CG = -3, -2; H(CG) = 0 if CG = -1, 0; H(CG) = -0.5 if CG = 1, 2 A learner needs to invest 50% more time in case of knowledge deficiency (CG=-2) and 50% less time in case of a knowledge surplus (CG=1 or 2). The required study time is in line with the LA estimated study time if the learner has adequate prior knowledge (CG=-1, or 0). Success = Scaled Effort + H(CG) - CG = Competence GAP; - Scaled Effort: 1, 2, 3 - H(CG): H(-2) = -2, H(-1) = 0, H(0) = 0, H(1) = 0, H(2) = 0 → output [-2,0] The learner successfully completes a LA if Success = 0, otherwise the LA examination fails.	Studying LA
- Success	Learner passes or fails LA-examination		Effort, LCL

Table 8.10 Graduation, satisfaction, time to graduate for a low-end goal (Step 2)

# Sub domains	Variables	Treatment						No treatment Control (C)	
		Ontology (O)		Peers (P)		OP		<i>M</i>	<i>SD</i>
3	Graduates	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
	– perc.	91.1	2.4	65.8	5.1	84.5	1.7	17.4	3.9
	– satisf.	39.2	4.0	6.6	2.0	23.1	3.1	2.5	1.3
	– time	1670	373	1843	352	1661	341	2286	413

Table 8.11 Graduation, satisfaction, time to graduate for a high-end goal (Step 2)

# Sub domains	Variables	Treatment						No treatment Control (C)	
		Ontology (O)		Peers (P)		OP		<i>M</i>	<i>SD</i>
3	Graduates	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
	– perc.	88.9	2.2	51.4	3.8	79.9	2.9	8.1	2.3
	– satisf.	87.2	2.3	53.7	2.4	63.8	4.2	38.0	12.1
	– time	5246	457	5576	587	5282	508	6061	536

Graduation. Analyses of variance showed – for all six conditions a significant difference in the percentage of Graduates between all groups. More detailed analysis showed that the control group always had a significant smaller percentage of graduates than all other treatment groups. This confirms H1(Part 1) that PRS recommendations yield to more graduation than no recommendations.

Satisfaction. Analyses of variance showed for all six conditions a significant difference in the percentage of maximum satisfaction at graduation between all groups. More detailed analysis, namely multiple comparisons all with Bonferroni’s correction, showed that the control group always had a significant smaller percentage of graduates with maximal satisfaction than any treatment group. This confirms our H1 (Part 2) that PRS recommendations yield to more satisfied graduation than no recommendations.

Time. Analyses of variance showed for all six conditions a significant difference in time to graduate between all groups. More detailed analysis, namely multiple comparisons all with Bonferroni’s correction, showed that the control group always needed significantly more time to graduate than any treatment group. This confirms H1 (Part 3) that PRS recommendations yield to faster graduation than no recommendations.

Hypothesis 2 (hybrid recommendations from PRS yield more, more satisfied, and faster graduation than purely CF-based RS) was also tested by using regular statistical methods, namely Bonferroni’s correction when using multiple comparisons at a more detailed level, using mean difference between two groups (P-group

and OP-group), $p < 0.05$. These multiple comparisons were always conducted after analyses of variance on the more global level. As before, results for all measurement variables are presented in Tables 8.10 and 8.11.

Graduation. Multiple comparisons, using Bonferroni's correction, revealed – for all six conditions – that the OP-group had significantly more graduates than the P-group, using mean difference between two groups, $p < 0.05$. This confirms Part 1 of H2.

Satisfaction. Multiple comparisons, using Bonferroni's correction, revealed – for all six conditions – that the OP-group had significantly more graduates with maximum satisfaction than the P-group, using mean difference between two groups, $p < 0.05$. This confirms Part 2 of H2.

Time. Multiple comparisons, using Bonferroni's correction, revealed – for all six conditions – the OP-group needing significantly less graduation time than the P-group, using mean difference between two groups, $p < 0.05$. This confirms Part 3 of H2.

Including more realism and flexibility at Step 2 has shown similar results as with Step 1, but now more realistic and also confirming the robustness of our conceptual simulation model. However, a possible drawback of the conceptual simulation models from previous Steps might be their rather intensive data-gathering aspect of the hybrid RS (reflected by 'OP'). Therefore, it was examined whether Ratings-based Hybrid RS could tackle this problem. In other words, in Step 3 the focus would be on practical implementation issues when designing a PRS-system.

8.2.3 Method Step 3: Benefits of Ratings-Based Hybrid RS

During this step, the main aim was to verify the effects of Ratings-based hybrid recommendation strategies in e-learning. In both previous steps, it was shown that hybrid recommendation strategies in e-learning were beneficial, but included an ontology-component that could still result in a rather heavy-weight PRS ('OP'). For this reason, it was examined whether rating-based recommendations could be a practically feasible alternative for ontology-based recommendations, as these might be expected to address the aim of more, faster, and more attractive graduation most closely. Our two main hypotheses for this step are:

H1: PRS recommendations yield more, more satisfied, and faster graduation than no recommendations

H2: ontology-based and rating-based recommendations from PRS show no differences for graduation, nor satisfaction, nor time to graduate.

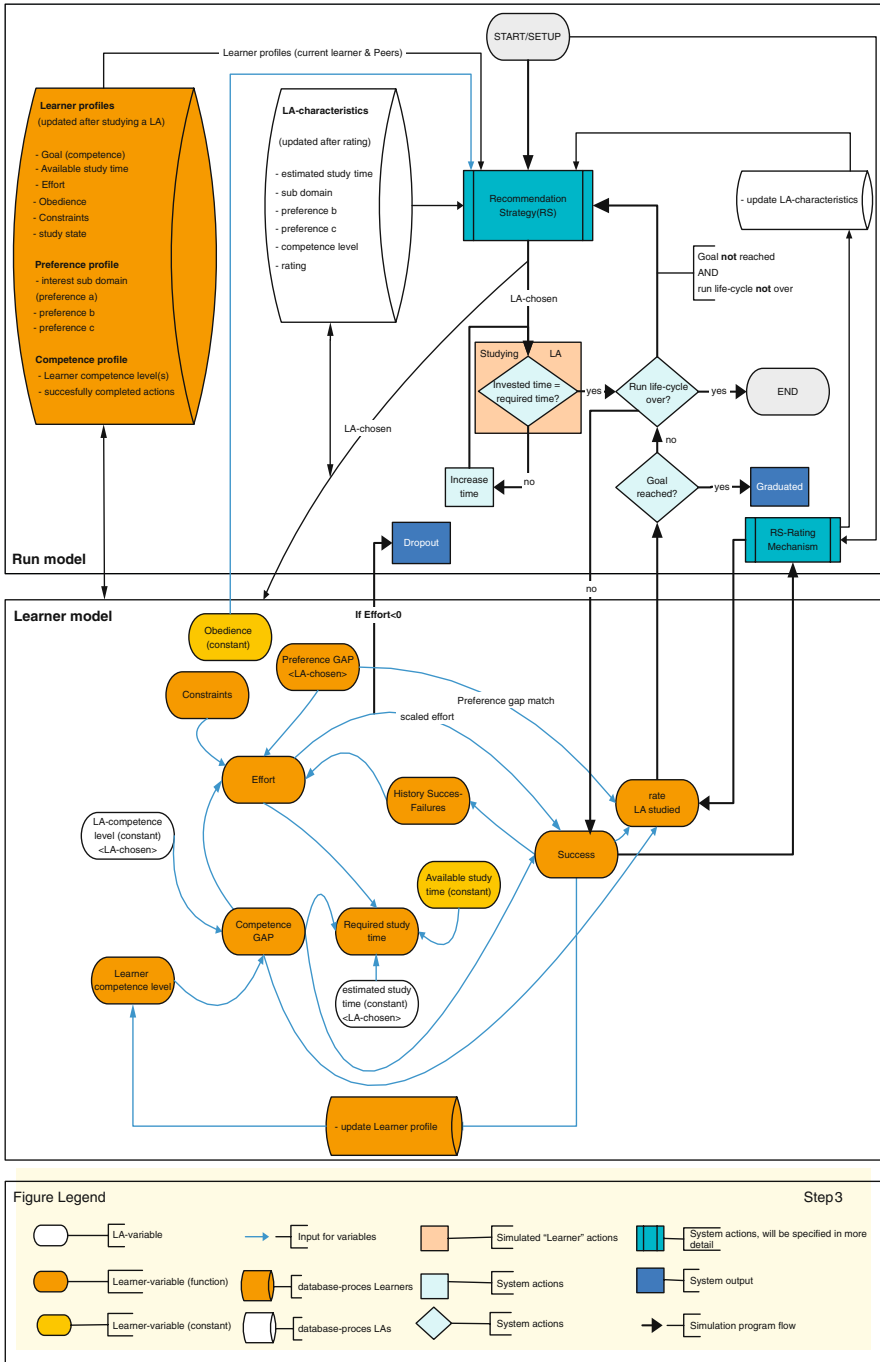


Fig. 8.5 Conceptual simulation model and simulation program flow (Step 3)

Table 8.12 Overview variables in conceptual simulation model Step 3 – different from Step 2 – and their implementation within the simulation

Variable	Description	Implementation in simulation		
		Range (initialisation)	Formula	Input for
Learner profile				
– All the same as in Step 2				
LA characteristics				
– All the same as in Step 2	–	–	–	–
AND	Perceived usefulness after studying	‘missing’, updated after each completion, integer [1,5]	Yes	RS
– Rating				
– Preference GAP	Alignment preferences and flavours	Integer [-3, 0]	Yes	Effort
– Preference gapmatch	Alignment preferences and flavours	Integer [-1,0], more coarse as ‘preference GAP’ [-3,0] → [-1,0]		Rate LA studied
– Competence GAP	Alignment LCL and CL of LA	Integer [-2, 2]	Yes	Effort, Success, rst, rate LA studied
– Required study time (rst)	Invested time before LA-examination	[50, 100, 150]	Yes	Studying LA
– Success	Learner passes or fails LA-examination	Boolean	Yes	Effort, LCL, rate LA studied

Figure 8.5 shows an integrated picture of the program flow in the simulation environment (the thick lines and arrows) and the conceptual simulation model for Step 3. Table 8.12 presents an overview of variables in the conceptual simulation model and their implementation within the simulation. Table 8.13 summarises relations and formulas, describing pedagogical rules in the conceptual simulation model. These rules are used within the RS. Please note that the model at Step 3 only differs from the model at Step 2 by inclusion of rating. For this reason, Tables 8.12 and 8.13 only include differences and will refer to Tables 8.8 and 8.9 if no differences with the Step 2 model are present. The model is implemented in such a way that weighting values for all variables could be easily adjusted. This is very helpful if more empirical data would become available.

The measurement variables are exactly the same as in the previous step. Application of the RS ultimately results in a LA-chosen by the learner (Fig. 8.6). The simulation can deal with seven treatment groups and one control group. The RS can partly be defined in the simulation setup. Recommendations can be based upon: Ontology (O), Peers (P), Ontology and Peers (OP), resulting in three treatment groups. Rating can be switched on or off in the setup thereby enabling four additional treatment groups: Ratings (R), Ontology and Ratings (OR), Peers and ratings

Table 8.13 Formulas and description for variables (different from Step 2) with changing values in the conceptual simulation model for Step 3

Variable	Description	Formula in simulation	Input for
- Learner competence level	Goal accomplishment	Exactly the same as for Step 1	Competence GAP
- Successfully completed LAs	Contribution towards goal		
- Effort	Investment to study and satisfaction during study	Exactly the same as for Step 2	Success, dropout, required study time
- Rating (updated each time this LA is studied)	Perceived usefulness after studying	Rating = $[w_3 * \text{previous LA-rating} + w_1 * G \text{ (individual rating)} + w_2 * H(\text{CF-rating})] / [w_1 + w_2 + w_3]$ - $w_1 = 0.25, w_2 = 0.25, w_3 = 0.5; w_1 + w_2 + w_3 = 1$ - previous LA-rating = average rating for all learners having studied this LA so far - CF-rating = average rating for ad-hoc group-members to which the current learner belongs when completing this LA (using Slope One Algorithm by Lemire and Maclachlan (2005)). - individual rating = $[J(\text{Success}) + K(\text{Preference gapmatch}) + L(\text{Competence GAP})] / [10]$ - $J(0)=25, J(1)=35$ - $K(-1)=-7.5, K(0)=7.5$ - $L(-2)=-7.5, L(-1)=0, L(0)=L(1)=7.5, L(2)=-7.5$ - individual rating (i.e., Rate LA, studied): 1, 2, 3, 4, 5 (Note: 5 for successfully completed, preference GAP is 0, and CG=0 or 1) - $H(\text{CF-rating}) \rightarrow 1, 2, 3, 4, 5$ (CF-rating (average) should be round off)	RS
- Preference GAP	Alignment preferences and flavours	Similar as for Step 1, however here three preference variables are used	Effort
- Preference gapmatch	See Step 2	Exactly the same as for Step 2	Rate LA studied
- Competence GAP	See Step 2	Exactly the same as for Step 2	Effort, Success, rst, rate LA studied
- Required study time (rst)	Learner passes or fails LA-examination	Exactly the same as for Step 2	Studying LA
- Success			Effort, LCL, rate LA studied

(PR), Ontology, Peers and Ratings (OPR). Learners choose the highest rated LA matching their interest.

Figure 8.1 also shows the simulation program flow. During the setup, exactly the same characteristics as for Step 2 can be considered/defined, although it is evident that now can be chosen from a larger variety of possible Recommendation Strategies and the results will differ (Fig. 8.3). The setup procedure initialises the environment in a similar way as for Step 1.

As for previous steps, every condition in the simulation was replicated 12 times (i.e., $N = 12$ runs). The simulation outcomes can be found at: <http://hdl.handle.net/1820/1491>. Each condition included seven treatment groups (O, P, R, OP, OR, PR, and OPR) and one control group (C). The same six conditions as before were included: goal (2) \times sub domain (3). In order to save space, we will focus on the groups (O, P, R, C) and on two conditions, namely a low-end goal for 3 sub domains and a high-end goal for 3 sub domains.

[Explanation Fig. 8.5: see also Explanation Fig. 8.6, Tables 8.8, 8.9, 8.12, 8.13 and Fig. 8.6]

Except for [rating], all LA-characteristics remain unchanged. Learners' rating of a LA ([rate LA-studied]) is influenced by whether or not the learner successfully completes this LA, [Preference gapmatch], and [Competence GAP].

Results Step 3

Hypothesis 1 (PRS recommendations yield more, more satisfied, and faster graduation than no recommendations) was tested by using regular statistical methods, such as analyses of variance on a global level, continuing with Bonferroni's correction when using multiple comparisons at a more detailed level. Here, multiple comparisons were always conducted between the control group and one of the seven treatment groups. Data means and standard deviations for all measurement variables are presented in Tables 8.14 and 8.15.

Graduation. Analyses of variance showed – for all six conditions a significant difference in the percentage of Graduates between all groups. More detailed analysis showed that the control group always had a significant smaller percentage of graduates than all other treatment groups. This confirms H1 (Part 1) that PRS recommendations yield to more graduation than no recommendations.

Satisfaction. Analyses of variance showed for all six conditions a significant difference in the percentage of maximum satisfaction at graduation between all groups. More detailed analysis, namely multiple comparisons all with Bonferroni's correction, showed that the control group always had a significant smaller percentage of graduates with maximal satisfaction than any

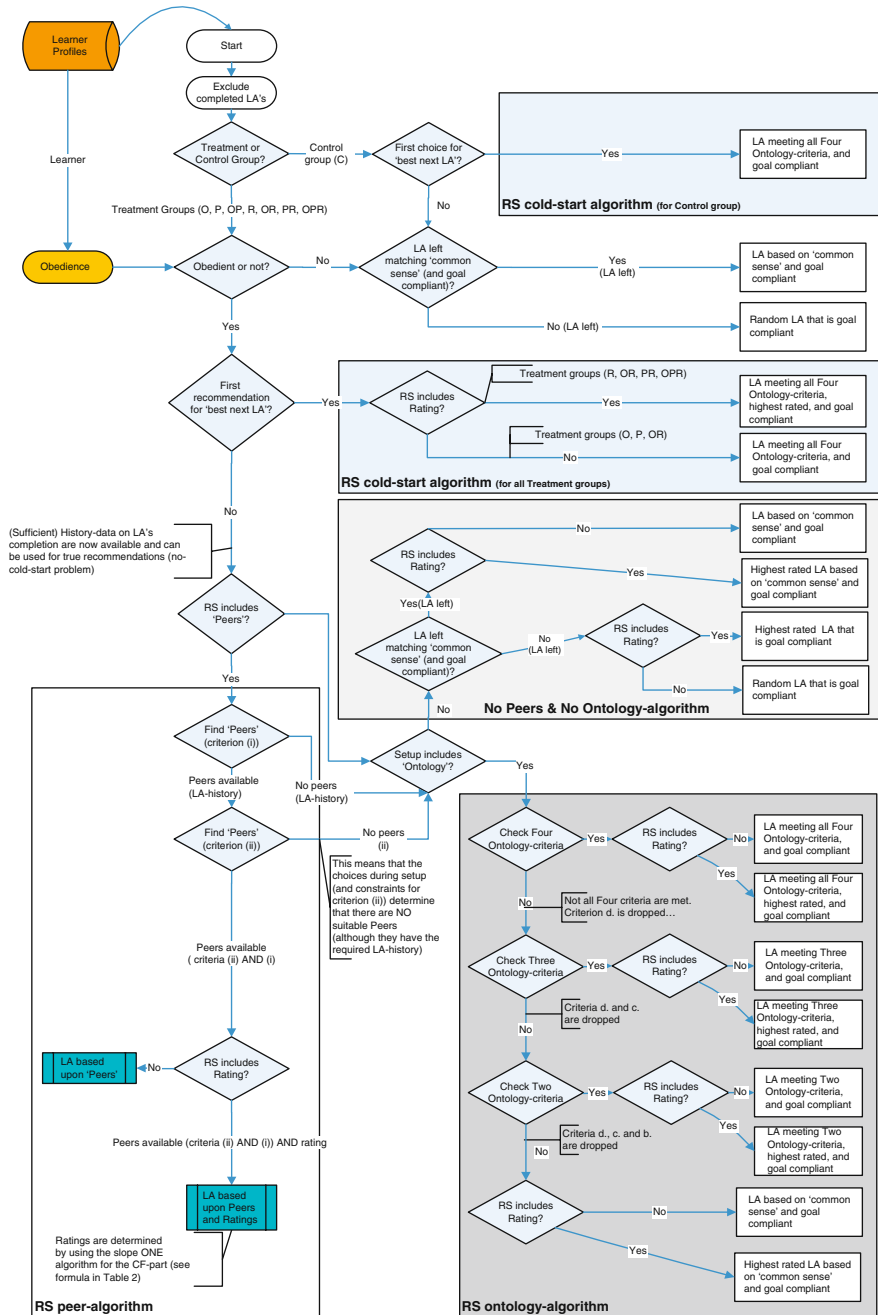


Fig. 8.6 Recommendation strategy (RS) for the simulation program at Step 3 in detail

treatment group. This confirms our H1 (Part 2) that PRS recommendations yield to more satisfied graduation than no recommendations.

Time. Analyses of variance showed for all six conditions a significant difference in time to graduate between all groups. More detailed analysis, namely multiple comparisons all with Bonferroni's correction, showed that the control group always needed significantly more time to graduate than any treatment group. This confirms H1 (Part 3) that PRS recommendations yield to faster graduation than no recommendations.

Hypothesis 2 (ontology-based and rating-based recommendations from PRS show no differences for graduation, nor satisfaction, nor time to graduate) was also tested by using regular statistical methods, namely Bonferroni's correction when using multiple comparisons at a more detailed level, using mean difference between two groups (O-group and R-group), $p < 0.05$. These multiple comparisons were always conducted after analyses of variance on the more global level. As before, results for all measurement variables are presented in Tables 8.14 and 8.15.

Graduation. For the low-level goal, multiple comparisons, using Bonferroni's correction, revealed no differences in the percentage of Graduates between the O-group and the R-group, using mean difference between two groups, $p < 0.05$. For the high-level goal, the same statistical methods revealed that the O-group had significantly more graduates if the number of sub domains were three or four. For two sub domains with the high-level goal, both groups showed no differences in the percentage of Graduates. This partly confirms Part 1 of H2.

Satisfaction. Multiple comparisons, using Bonferroni's correction, revealed significant differences in percentage of Graduates with maximum satisfaction between the O-group and the R-group, using mean difference between two groups, $p < 0.05$. For all 6 conditions, the O-group had significantly more satisfaction at graduation than the R-group. This rejects Part 2 of H2.

Time. Multiple comparisons, using Bonferroni's correction, showed some significant differences in graduation time between the O-group and the R-group, using mean difference between two groups, $p < 0.05$. The R-group consistently needed less graduation time than the O-group for the low-level goal. However, the R-group needed more graduation time than the O-group for the high-level goal with three or four sub domains. There were no differences between both groups for the high-level goal with two sub domains. This partly confirms Part 3 of H2.

Finally, it was found that for all RS an increasing number of sub domains reduce the effect of PRS on all three measures. Some additional simulations in which the number of sub domains was further increased showed a similar, but quite smoothly trend.

Table 8.14 Graduation, satisfaction, time to graduate for a low-end goal (Step 3)

# Sub domains	Variables	Treatment						No treatment Control (C)	
		Ontology (O)		Peers (P)		OP			
		<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
3	Graduates								
	– perc.	91.1	2.4	65.8	5.1	93.9	1.3	17.4	3.9
	– satisf.	39.2	4.0	6.6	2.0	33.3	3.1	2.5	1.3
	– time	1670	373	1843	352	1661	341	2286	413

Table 8.15 Graduation, satisfaction, time to graduate for a high-end goal (Step 3)

# Sub domains	Variables	Treatment						No treatment Control (C)	
		Ontology (O)		Peers (P)		OP			
		<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
3	Graduates								
	– perc.	88.9	2.2	51.4	3.8	83.8	2.6	8.1	2.3
	– satisf.	87.2	2.3	53.7	2.4	61.8	3.0	38.0	12.1
	– time	5246	457	5576	587	5326	430	6061	536

8.3 Conclusion

The description and application of aforementioned three Steps demonstrates that a stepwise approach towards simulations supported us in the articulation of light-weight PRS requirements. Our study started with a focussed expert discussion in which variables, their relationships, and associated literature with existing research were identified as input for all steps. In the first step, focus was given towards simplicity. The second step focused on more reality, whereas the third step focused on practical implementation issues. In fact, we followed a well established approach in social science (Gilbert and Troitzsch 1999), but hitherto sparsely used for recommendations in e-learning. This last section starts with a summary of results and discussion.

8.3.1 Important Findings and Discussion

This simulation study clearly demonstrates the advantage of using pedagogy-oriented PRS recommendations in e-learning. For all steps, it was shown that these yield to more, and more satisfied graduation than without recommendations.

This study also shows that – in particular for low-goal achievement – rating-based recommendations are a good light-weight alternative for intensive data maintenance approaches needed by ontology-based recommendations. For a low-level goal, rating-based recommendations yield similar graduation figures and even less time to graduate than ontology-based recommendations, although learners are less satisfied. For a high-level goal, ontology-based recommendations mostly result in

better graduation figures and less time to graduate than rating-based recommendations, again learners are less satisfied with rating-based recommendations. At a practical level, differences in time to graduate at both levels might probably be ignored as they are within the range of a couple of percent. For similar reasons, differences in graduation figures for the high-level goal might probably be ignored.

This simulation study shows that rating-based recommendations in comparison with peer-based recommendations result in more, more satisfied, and faster goal achievement. Both RS use a similar mechanism, but whereas rating-based recommendations use 'peers with similar LA-history', peer-based recommendations additionally use 'peers with similar profile'. Rating-based recommendations are somewhat more light-weight than peer-based recommendations and thus preferred.

Finally, for all RS it was found that an increasing number of sub domains reduces the effect of PRS on all three measures. This seems reasonable as our conceptual simulation model only takes one preference sub domain into account. In addition, more LAs than the 400 included in our simulations might be needed for learners to be able to achieve the goal without having to study non-preferred LAs. In other words, the more sub domains the more chance that there are too few LAs within the preferred sub domain. Although an increasing number of sub domains reduces the effect of PRS on all three measures, it is interesting to see that this can be different for various treatments. For example for the level-1 goal, whereas the mean percentage graduates for P drops from 70.5% (2 sub domains) to 60.8% (4 sub domains), for R the mean percentage graduates drops from 94.5 to 89.4%. We have no explanation for this significant difference. It shows that R is more robust than P with respect to the number of sub domains. But, it also demonstrates that a conceptual simulation model cannot unequivocally predict its outcomes.

In another study, the relevance of inclusion of a cold-start algorithm in the RS was confirmed (Nadolski et al. 2009).

8.3.2 Practical Implications

Every condition in the simulation was replicated 12 times (i.e., $N = 12$ runs) to justify the use of classical statistic techniques on resulting data (see Law and Kelton 2000, p. 496). A possible drawback of this approach could be that the law of large numbers ($n = 3000$ learners for each group, i.e., 12 runs with 250 learners for each group) might explain occurrences of significant differences. However, as almost all effects were also prevalent at the level of separate runs, this possible explanation can be ruled out in our simulation study. Results on the level of separate runs also showed that the conceptual simulation model caters for randomness. Some randomness in behaviour is crucial in a LN for facilitating the discovery of new solutions (Koper 2005).

To conclude, this simulation study shows the prospective benefit of using ratings within a RS of the PRS. Although it is argued that a rating-based PRS results in a rel-

atively light-weight PRS, it could still put considerable demands on the needed technical infrastructure needed as information in LN is highly distributed. So, although in principle, this could work, but actual implementation is ‘the proof of the pudding’. Furthermore, Farzan and Brusilovsky (2006) report about the ‘users-do-not-like-to-rate’ phenomenon. This means that actual deployment of such a system should tackle this problem in a way similar to as was suggested by aforementioned authors, namely by providing a clear incentive/benefit to learners for providing such ratings, for example by receiving credits for doing so (Hummel et al. 2005). On the other hand, other studies mention possible drawback of providing clear incentives on willingness to share knowledge (Bock and Kim 2002). Finally, a considerable number of learners and learning actions are needed to make this approach work.

8.3.3 *Limitations*

Our simulation study has several limitations (see also Koper 2005). We will address four limitations.

The main one comes from the limited availability of real data. This boils down to the question whether learners really behave as is modelled in the conceptual simulation model. The conceptual simulation model is derived from various theories (see also Koper 2005), but at a more detailed level, we were sometimes forced to make arbitrary choices. For example: obedience towards recommendations is personalised, but is still supposed to be stable for each learner. As Koper (2005), we expect that obedience will be positively influenced by the usefulness of the previous recommendations (i.e., increasing confidence), but will also depend upon the way and format in which they are presented (see Farzan and Brusilovsky 2006). Indeed, further usability-studies are needed that make an inventory of what learners expect from recommendation systems as well as their – partly individual – preferences towards information presentation formats for such systems.

The second limitation stems from the fact that a conceptual model always represents part of the world, with as drawback that features could still be missing, notwithstanding the expert consultation that was carried out to verify the current model. Four examples are included to clarify this. First, it was not explicitly taken into account that learners’ choice in performing a recommended LA depends on the enrolment-costs of that LA. Although this could be represented by preference b or c , there might still be the need for other weighting factors than were used in the current model. Second, we did not take into account some ‘constraints’ that could act as ‘negative feedback’ in the conceptual simulation model. For example: LAs might have some enrolment constraints (maximum number of learners, limited start-time, a.s.o.). Third, learners do not interact with each other in the LN, only indirect social interaction is modelled. However, in reality there is always a combination of indirect and direct social interaction influencing the choices made. Fourth, learners can have preference for one sub domain only, whereas specific learners could

like several sub domains equally or would like to indicate their preferences more gradually.

The third limitation comes from the simplification towards competences. Our conceptual simulation model does not deal with competence-hierarchies. These are complex to describe and the PRS system should be able to trace how completion of a certain learning action can contribute to the – partial-attainment of more than one competence. The learning path specification which is currently under development could tackle this problem (Janssen et al. 2008).

The fourth and final limitation of this study is that it shows that somewhat ontology-based recommendation, namely a limited set of LA-characteristics and learner-characteristics, should preferably be part of the hybrid approach. However, this can induce considerable workload that could be minimised when using a limited vocabulary of meta data. But, how to identify and maintain such limited vocabulary that can be expected to be partly domain dependent?

8.3.4 Future Plans

Our research focuses on light-weight system-based personalised recommendations within LN. It is debatable whether this approach is sufficient or should be augmented with human-based recommendations as is suggested by QSIA (Rafaeli et al. 2004, 2005). As community driven information (CF) becomes more important within the recommendation strategy, it should be noted that this has a bias to popular tastes. Learners with an unusual taste may get fewer qualitative recommendations, and others are unlikely to be recommended unpopular items (of high quality). Further research is needed to justify the augmentation with human-based recommendations as these can be very costly and practically almost impossible because of limited availability of experts.

The most important finding from our simulation study is the prospective benefit of ratings. Future field experiments on LNs with real learners should verify the value of light-weight system-based personalised recommendations. Such experiments can also assist the further articulation of a conceptual simulation model and subsequent simulation studies. Such iterations between simulations, system development, and field experiments will provide valuable insights into the technical infrastructure needed to provide sound personalised recommendations to professionals in Learning Networks.

References

- Andronico, A., Carbonaro, A., Casadei, G., Colazzo, L., Molinari, A., Ronchetti, M.: *Integrating a Multi-Agent Recommendation System into a Mobile Learning Management System*, 2003 [Retrieved October 18, 2007 from <http://w5.cs.uni-sb.de/~krueger/aims2003/camera-ready/carbonaro-4.pdf>]
- Berlanga, A.J., Van den Berg, B., Nadolski, R., Drachsler, H., Hummel, H.G.K., Koper, R.: Towards a model for navigation support in Learning Networks. *ePortfolio Confer-*

- ence, October 18–20, 2007, Maastricht, The Netherlands [Retrieved October 5, 2008 from <http://hdl.handle.net/1820/1015>]
- Bock, G.W., Kim, Y.: Breaking the myths of rewards: an exploratory study of attitudes about knowledge sharing. *Information Resources Management Journal* **15**(2), 14–21 (2002)
- Bolman, C., Tattersall, C., Waterink, W., Janssen, J., Van den Berg, B., Van Es, R., Koper, R.: Learners' evaluation of a navigation support tool in distance education. *Journal of Computer Assisted Learning* **23**, 384–392 (2007)
- Claypool, M., Gokhale, A., Miranda, T., Murkinov, P., Netes, D., Sartin, M.: Combining content-based and collaborative filters in an online newspaper. *Conference ACM SIGIR Workshop on Recommender Systems: Algorithms and Evaluation*, August 19, 1999, Berkeley, CA, USA
- Cranen, K.: Using digital consults and determinants in following up advices (Master thesis, Enschede, The Netherlands: Twente University 2007)
- Drachslar, H., Hummel, H.G.K., Koper, R.: Personal recommender systems for learners in lifelong learning: requirements, techniques and model. *International Journal of Learning Technology* **3**(4), 404–423 (2008a)
- Drachslar, H., Hummel, H.G.K., van den Berg, B., Waterink, W., Eshuis, J., Nadolski, R.J., Berlanga, A.J., Boers, N., Koper, R.: Effects of the ISIS recommender system for providing personalised advice to lifelong learners: an experimental study in the domain of Psychology. *4th Edumedia Conference 2008 Self-organised Learning in the Interactive Web – Changing Learning Culture?* June 2–3, 2008b, Salzburg, Austria: CEUR
- Farzan, R., Brusilovsky, P.: AnnotatEd: A Social Navigation and Annotation Service for Web-based Educational Resources. In T.C. Reeves, S.F. Yamashita (Eds.): *Proceedings of World Conference on E-Learning, E-Learn 2006*, pp. 2794–2802, October 13–17, 2006, AACE, Honolulu, HI, USA,
- Gilbert, N., Troitzsch, K.G.: *Simulation for the Social Scientist* (Buckingham, Philadelphia, USA: Open University Press 1999)
- Good, N., Schafer, J.B., Konstan, J.A., Borchers, A., Sarwar, B., Herlocker, J., Riedl, J.: Combining collaborative filtering with personal agents for better recommendations. *Proceedings of AAAI*, pp. 439–446, 1999
- Hummel, H.G.K., Burgos, D., Tattersall, C., Brouns, F.M.R., Kurvers, H.J., Koper, E.J.R.: Encouraging contributions in learning networks by using incentive mechanisms. *Journal of Computer Assisted Learning*, **21**(5), 355–365 (2005).
- Hummel, H.G.K., Van den Berg, B., Berlanga, A.J., Drachslar, H., Janssen, J., Nadolski, R., Koper, R.: Combining social- and information-based approaches for personalised recommendation on sequencing learning activities. *International Journal of Learning Technology* **3**(2), 152–168 (2007)
- Janssen, J., Berlanga, A.J., Vogten, H., Koper, R.: Towards a learning path specification. *International Journal of Continuing Engineering Education and Lifelong Learning* **18**(1), 77–97 (2008)
- Konstan, J., Miller, B., Maltz, D., Herlocker, J., Gordon, L., Riedl, J.: GroupLens: Applying collaborative filtering to usenet news. *Communications of the ACM* **40**(3), 77–87 (1997)
- Koper, R.: Increasing learner retention in a simulated learning network using indirect social interaction. *Journal of Artificial Societies and Social Simulation* **8**(2) (2005) [Retrieved October 30, 2007 from <http://www.jasss.soc.surrey.ac.uk/8/2/5.html>]
- Koper, R., Olivier, B.: Representing the learning design of units of learning. *Educational Technology & Society* **7**(3), 97–111 (2004)
- Law, A.M., Kelton, W.D.: *Simulation Modeling and Analysis* (Boston: McGrawHill 2000)
- Lemire, D., MacLachlan, A.: Slope One Predictors for online rating-based collaborative filtering. *SIAM Data Mining (SDM'05)*, April 21–23, 2005, Newport Beach, California
- McCalla, G.: The ecological approach to the design of e-Learning environments: Purpose-based capture and use of information about learners. *Journal of Interactive Media in Education* (2004) [Retrieved October 18, 2007 from <http://www-jime.open.ac.uk/2004/7>]
- Melville, P., Mooney, R.J., Nagarajan, R.: Content-boosted collaborative filtering for improved recommendations. *Proceedings of 18th National Conference on Artificial Intelligence*, pp. 187–192, July 28–August 1, 2002, Edmonton, Alberta, Canada

- Nadolski, R.J., Van den Berg, B., Berlanga, A.J., Drachsler, H., Hummel, H.G.K., Koper, E.J.R.: Simulating light-weight personalised recommender systems in learning networks: A case for pedagogy-oriented and rating-based hybrid recommendation strategies. *Journal of Artificial Societies and Social Simulation (JASSS)*, **12**(1), 4 (2009)
- Pazzani, M.J.: A framework for collaborative, content-based and demographic filtering. *Artificial Intelligence Review* **13**(5), 393–408 (1999)
- Rafaeli, S., Barak, M., Dan-Gur, Y., Toch, E.: QSIA: a Web-based environment for learning, assessing and knowledge sharing in communities. *Computers & Education* **43**, 273–289 (2004)
- Rafaeli, S., Dan-Gur, Y., Barak, M.: Social recommender systems: recommendations in support of e-learning. *Journal of Distance Education Technologies* **3**(2), 29–45 (2005)
- Ryan, R.M., Deci, E.L.: Self-determination theory and the facilitation of intrinsic motivation, social developments and well being. *American Psychologist* **55**, 68–78 (2002)
- Soboro, I.M., Nicholas, C.K.: Combining content and collaboration in text filtering. *Proceedings of the IJCAI Workshop on Machine Learning in Information Filtering*, pp. 86–91, August 1999, Stockholm
- Tang, T.Y., McCalla, G.: Beyond learners' interest: personalized paper recommendation based on their pedagogical features for an e-learning system. *Proceedings of the 8th Pacific Rim International Conference on Artificial Intelligence (PRICAI 2004)*, August 9–13, 2004a, Auckland, New Zealand
- Tang, T.Y., McCalla, G.: Utilizing artificial learners to help overcome the cold-start problem in a pedagogical-oriented paper recommendation system. *Proceedings of AH 2004: International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, August 23–26, 2004b, Eindhoven, The Netherlands
- Van Setten, M.: *Supporting People in Finding Information: Hybrid Recommender Systems and Goal-Based Structuring* (Doctoral thesis, Enschede, The Netherlands: Telematica Instituut 2005)
- Vygotsky, L.: *Mind in Society: The Development of Higher Psychological Functions* (Cambridge: Harvard University Press 1978)
- Walker, A., Recker, M.M., Lawless, K., Wiley, D.: Collaborative information filtering: a review and an educational application. *International Journal of Artificial Intelligence in Education* **14**, 1–26 (2004)
- Wilensky, U.: NetLogo. <http://ccl.northwestern.edu/netlogo/>. *Center for Connected Learning and Computer-Based Modeling*, (1999), Evanston, IL: Northwestern University

Concerns about a lack of student guidance and support in approaches that put learners centre stage and in control of the learning process are understandable, especially in dynamic environments that offer a considerable number of options to choose from. Studies relating to navigation in hypermedia systems show that especially novices in a field need to be provided with appropriate content structure and navigational support to reduce disorientation and help them develop a structural representation of the knowledge domain (Chen et al. 2006).

Darken and Silbert (1993) distinguish three dimensions in their classification of virtual environments – *size*, *density* and *activity* – that can be used to characterise Learning Networks as well. Learning Networks are likely to be large, to the extent that it is not possible to get a view of the entire network in detail. They can also be dense or even cluttered: if there are several ways to achieve the same learning goals, ‘learning goal’ will be less powerful as a cue for the selection of a learning path. Other cues or criteria will have to be taken into account as well, like for instance costs or study load. Finally, Learning Networks are likely to show a high level of activity. They are highly dynamic: not only will learning paths and activities change due to updates but they might be removed as well and new activities and paths will be added. This makes the process of finding one’s way in a Learning Network rather complicated. Though the way finding problem is not unique to Learning Networks, but part of educational practice today (Martinez and Munday 1998; Simpson 2004; Yorke 1999), the mere size, density and activity of Learning Networks demands that a generic solution is found.

For the sake of clarity we begin by stating that the notion of ‘finding one’s way’ in a Learning Network, can be considered in a quite broad sense so that it not only involves finding suitable learning paths but also more general orientation and goal definition, getting acquainted with policies and practices, getting to know other people in the network, etc. The concept of *navigation* refers to: ‘the process by which people control their movement using environmental cues and artificial aids such as maps so that they can achieve their goals without getting lost’ (Darken and Silbert 1993, p. 157). Though knowing policies, practices and people in the Learning Network can considerably contribute to this process of goal achievement, we will use the term navigation here in a more restricted sense indicating the process of selecting a learning path suitable to achieve one’s learning goals and the ensuing process of proceeding along this learning path.

A possible way of providing navigational support is to recommend a best next step along a learning path, based on the way other learners proceeded (Brooks et al. 2005; Drachsler et al. 2008; Janssen et al. 2007). By means of social filtering techniques navigational clues can be inferred from the way other learners with comparable goals and preferences navigated a learning path. Determining a path to follow however means that all kinds of possible restrictions must be taken into account. For instance, it might not make much sense for a learner to start developing competence Y unless competence X has been acquired. It might even constitute a formal requirement within a curriculum. This requires that learning paths are described in a way that makes this type of information available and, preferably, amenable to automated processing. Apart from the issues of how to support navi-

gation of a particular learning path and how to clarify possibly related constraints, there is the prior concern of how to support learners in choosing a learning path from the available options. In order to enable description, comparison and selection of suitable learning paths, a formal and uniform way to describe learning paths has been developed (Janssen et al. 2008). This *learning path specification* captures the characteristics of a learning path in a formal way to offer learners, teachers, administrators, systems and applications like the Personal Competence Manager (Vogten et al. 2008) the information they need to identify suitable learning paths and to navigate them. The main aim is to offer learners a ‘viewer-centred’ overview (Lueg and Bidwell 2004), first of suitable learning paths they can select from and then of suitable ways to proceed along the learning path that has been selected. This chapter further illustrates how this works: how the learning path specification can be used to realise this, starting from a scenario of a novice learner who wants to know more about interior design.

9.2 Scenario

Currently, a learner who is about to engage in a new field of professional competence development is likely to ask around, have a look at different educational providers to see what they offer, for instance through a search on Internet. Entering ‘interior design course’ for a search in Google, currently will result in over 58 thousand hits, referring to all kinds of interior design courses and pages referring to this type of courses, at different levels, some accredited others not, with different price tags attached, with varying study load, etcetera. Apart from the fact that this clearly represents a case of information overload, a novice in the field has no idea where to start, even if some course titles might offer a hint: ‘introduction’, ‘basics’. Otherwise she has no idea how the domain is ‘organised’: the competences it involves and how they are related. Our learner is likely to feel impeded being confronted with all the options and the variables to take into account comparing these options. Then she hears about a Learning Network on interior design and decides to try her luck there. Upon enrolment she is asked to specify a learning goal. ‘Wanting to know more about interior design’ is not sufficient: Learning Networks are envisaged to operate on the basis of competence maps, specifying competences and competence profiles to be acquired in a certain domain (Kickmeier-Rust et al. 2006). These competence maps represent competence levels and hierarchical relations between competences. So a learner may have to ‘translate’ somewhat general learning goals to specific competences she would like to attain. She might be helped in the process for instance by presentations of people with different competence levels and descriptions of what they are able to do. The learner could choose one of these ‘role models’: I would like to develop myself so that I can do what this person is doing. Or she might be matched to more advanced learners in the Learning Network who can guide her in defining desired competences. In any event some counselling or advice is likely to be needed at this stage of orientation and goal definition. Another issue at this same stage is to establish what competences and competence levels a learner has already

attained. For the purpose of this chapter it suffices to note that a learner profile will be created, which will describe the competences a learner has attained. This profile will be regularly updated, so that the moment a learner chooses a new goal and related path, it can be established whether the learner already has acquired some competence(s) that are part of this new path and hence can be considered ‘done’.

Let’s say our learner wants to know more about interior design because she works in a shop selling paints and wallpaper and would like to be better able to offer advice to customers. Let’s further assume that according to the competence map used in this Learning Network this competence is called *Offering interior design advice* which has three levels of proficiency: a basic, an advanced, and an expert level. The competence of offering interior design advice at a basic level implies a diversity of knowledge and skills such as *choosing colours*, *materials characteristics and applications*, *distinguishing style and design*, and *effective visual communication* (Fig. 9.1).

Figure 9.1 illustrates how the competences *choosing colours* and *materials characteristics and applications* are considered prerequisite to developing the competence *distinguishing style and design*. These competences can be attained through related learning paths and learning activities.

The learner in our example wants to cover the whole array of the competence at a basic level, so this is the learning goal she selects, rather than one of the sub-competences. The next stage for the learner is to find a learning path that will help her attain the desired competence at the desired level. At this point the learning path specification comes into the picture. As stated in the introduction the learning path specification defines characteristics of a learning path that offer learners the information they need to find suitable learning paths.

These characteristics are described in *elements* of the specification relating to the contents of the learning path (learning activities and learning goals) and to the structure (organisation, workflow) of the learning path. The structural aspects (organisation, workflow) offer the navigational clues, for instance concerning the

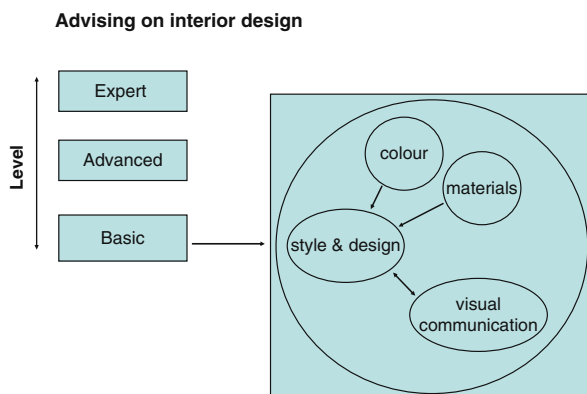


Fig. 9.1 Schematic representation of the competence *Offering interior design advice*

order in which to perform the learning activities, that are needed to support learners to *follow a particular learning path*. The next section will focus on characteristics identified by the learning path specification that are key to finding and following a learning path.

9.3 Key Elements to Describe Learning Paths

The learning path specification contains a number of elements to express contents and structure of a learning path. Requirements for the specification have been formulated based upon a review of literature on curriculum design and an analysis of different approaches to support comparison and/or selection of courses and programmes (Janssen et al. 2008). The same study revealed that we can draw on the existing IMS Learning Design specification (IMS-LD 2003) to describe learning paths. Table 9.1 gives an overview of the key elements of the learning path specification and indicates which elements are considered important to support learners in finding and following a learning path. Learning objectives, prerequisites and metadata are key elements when it comes to finding a suitable learning path. Metadata can be defined as data that provide information about data. In the context of a learning path: data that provide information about learning paths, e.g. title, language, studyload, costs. To describe these metadata the learning path specification refers to the IEEE LOM specification (IEEE/LOM 2002) using extensions where necessary (Duval et al. 2002).

Learning objectives and prerequisites are defined both at the learning path level and the level of its constituent parts: learning activities. This makes it possible to identify parts of a learning path a learner may be able to skip, given her profile and the competences she already has. For instance if the learner in our example had already acquired the competence regarding materials she could skip any learning activities related to this competence. This can be expressed using the conditions element of the learning path specification by stating that an activity can be considered completed if the learner profile indicates the competence level it refers to already has been achieved.

Learning activities may be clustered or ‘organised’ in a structure for a number of reasons e.g. the learner may have to perform one learning activity first in order to acquire the necessary prerequisite knowledge for a next learning activity (sequence) or the learner may be able to choose one or more learning activities from a defined subset (selection). These criteria might vary depending on the circumstances or learners’ backgrounds. This again could be expressed by specifying conditions or by specifying different learning paths.

Figure 9.2 explains how the structure element (selections and sequences) is used to model a learning path. The figure pictures a learning path leading to the basic level competence of *offering interior design advice* as represented in Fig. 9.1. According to this particular learning path our learner can achieve the competences *choosing colours, materials characteristics and applications*, and *distinguishing style and*

Table 9.1 Key elements of the learning path specification

Element	Explanation
Learning path	This is the container element of one or more learning activities that constitute a learning path. It has a title and an id so that the path (e.g. a professional competence development programme, course) can be presented and linked to. The learning paths goal and required entry levels are described through the constituent elements 'learning objectives' and 'prerequisites' (see below).
Learning activity	This element is used to describe each learning activity within the learning path, through the constituent elements 'learning objectives' and 'prerequisites' (see below).
Learning objectives	Learning objectives are specified at the level of each single activity as well as the aggregate level of a learning path. The element is used to define the learning goal by referring to competences.
Prerequisites	Like the learning objectives this element, which specifies the entry-requirements in terms of competences, is specified at the level of single learning activities and the aggregate learning path level.
Structure	This element is used to group learning activities that are somehow related, for instance because they compose a set a learner can choose from, or because they have to be studied in a particular order.
Method	The method element is used to order the defined learning activities and structures into a scenario that constitutes the 'learning flow' of the learning path.
Conditions	Conditions are used to further model the learning flow for instance to accommodate personalisation. So a condition could be: 'if the learner has chosen to complete activity / course X rather than Y, then course Z should be delivered in the mode where examples are given referring to course X'.
Metadata	The specification includes a set of metadata which, apart from the above elements, are of importance to identify learning paths most in line with learner's needs. They provide additional search criteria: title, language, recognition, start conditions, study load, assessment, delivery mode, teaching place, delivery mode, start date, end date, contact hours, guidance, costs, and contact.

design each through a single learning activity. However, in order to acquire the competence 'effective visual communication' she has to perform a series of learning activities (sequence). As Fig. 9.1 shows, the knowledge and skills regarding colour and materials have to be mastered before the learner can proceed with style and design. Figure 9.2 reflects this by clustering the learning activities for colour and materials in a selection and by incorporating this cluster in a sequence together with the learning activity for style and design, to indicate that the activities for colour and materials have to be performed prior to the learning activity for the development of skills regarding style and design. The learning activities in a selection can be performed in random order. Additionally it is possible to specify a choice by indicating the number of learning activities of the total set the learner has to complete.

The sequence of four learning activities leading to the visual communication competence and the sequence leading to the attainment of the other three compe-

Table 9.2 Key elements of the learning path offering interior design advice

Element		Description
Learning Path	Title	Offering interior design advice within a year
	Learning objectives/Competence	<i>Offering interior design advice</i>
	Prerequisites	None
	Recognition	Yes
	Learning activity	
	Title	Picking and mixing colours
	Learning objectives	<i>Choosing colours</i>
	Prerequisites	None
	Learning activity	
	Title	Know your materials
	Learning objectives	<i>Materials characteristics and applications</i>
	Prerequisites	None
	Learning activity	
	Title	Style and design
	Learning objectives	<i>Distinguishing style and design</i>
Prerequisites	+ <i>Choosing Colours</i> + <i>Materials characteristics and applications</i>	
Learning activity	learning activity 1 of sequence of activities leading to competence Visual Communication etc.	
Structure	Selection (top level) of: Sequence:	1. (Selection: Choosing Colours / Know your Materials) 2. Style and Design
	Sequence:	Visual Communication 1, 2, etc.
Method Conditions		Present top level selection
		If learner portfolio indicates that materials competence already has been attained set learning activity ‘Know your materials’ to ‘completed’. If learner portfolio indicates a preference for written text over audio present learning activity ‘Style and Design’ in modus T (textual).

tences are unrelated, so the outer circle in Fig. 9.2 again represents a selection (Note that elaborate designs like this example are the result of a learning design process carried out by an educational provider.).

For the learner this means she can start with either the learning activity *materials*, or the learning activity *colour* or the first activity of the sequence *visual communication*, as will be further illustrated in Sect. 9.4.2.

To illustrate this example further, Table 9.2 shows how the key elements of the learning path specification (presented in Table 9.1) should be described.

Having explained the main elements of the learning path specification we will now return to our learner to have a closer look at the way she can be supported in finding and following a learning path, using the specification.




Fig. 9.2 Example of a learning path composition

9.4 The Learning Path Specification in Practice

9.4.1 Finding a Learning Path

Let's see what happens when our learner selects a goal (competence) that can be developed through a large number of different learning paths: i.e. there are a large number of learning paths available through the Learning Network which refer to this competence in the learning objectives element. The learner must be offered means to efficiently choose the learning path that best fits her needs. Taking a decision support perspective we distinguish two stages in the decision making process: screening and choice (Beach 1997; Rundle-Thiele et al. 2005). Screening involves selecting a number of options one wants to take into consideration. It means narrowing down the number of choice options to a number that can be 'managed'. Time and attention are scarce resources (Goldhaber 1997; Schwartz 2004) and our learner would rather invest these resources in developing competences than in comparing all the different ways in which this can be done. What is needed then, is a tool for the learner to select a limited set of learning paths based on additional criteria so that eventually perhaps three or four options remain to take a closer look at and finally choose from.

It makes sense to enable the learner to set possible additional criteria already at the stage where she indicates which competence she wants to develop. The overview of possible learning paths in the Learning Network would then offer a selection already taking into account these additional criteria. In fact quite a number of web portals in the field of education operate this way. Figure 9.3 offers an

example of a search interface used at the UK Learndirect website (<http://www.learndirectadvice.net/findacourse/>).

Of course the question of which criteria to include in advanced searches is open to debate. Since a learner in a Learning Network will be asked to specify a goal in relation to a competence map, subject keywords are less relevant to include in this context: learning paths will be presented based on the selected goal/competence,

Subject keyword(s) ^{*}
What do you want to learn?

Location
Where do you want to learn?
 from

or

Study type
Which type of courses are you looking for? [NB if you choose to include 'Self Study' it refers to distance learning - ie online or by mail and so the provider may not necessarily be in your local area.]

<input checked="" type="checkbox"/> Full Time	<input checked="" type="checkbox"/> Part Time Evening
<input checked="" type="checkbox"/> Part Time Day	<input checked="" type="checkbox"/> Short
<input checked="" type="checkbox"/> Weekend	<input checked="" type="checkbox"/> Self Study
<input checked="" type="checkbox"/> Customised	<input checked="" type="checkbox"/> Day/Block Release

Duration
How long would you like the course to last for?

Start date
When would you like to start the course?
 Include flexible start dates

Qualification
Which qualification (if any) are you interested in?

Fig. 9.3 Advanced search example from the Learndirect website <http://www.learndirectadvice.net/findacourse/>

rather than a subject keyword entry. In the learning path specification additional criteria are included as metadata: costs, study load, teaching place, guidance, assessment etc.

Some criteria will be especially relevant to the screening process, because they represent possibly ‘hard’ constraints on the part of the learner: language, costs, start date, study load, recognition, and delivery mode. If our learner has a maximum amount of time and/or money to spend and does not master any foreign languages, learning paths which do not fit these criteria must be filtered out in the screening process.

Other criteria are less clear-cut and likely to be considered in the choice stage – the more in depth comparison of a limited number of learning paths. This in depth comparison is likely to involve elements like description, guidance, and additional information that might be provided through a link.

There are quite a number of criteria that could be relevant to finding the most suitable learning path and not all criteria are equally relevant to all learners. Progressive disclosure of functionality could therefore contribute to help the learner focus on those criteria that are most relevant for her (Turbek 2008). Progressive disclosure is ‘a strategy for managing information complexity in which only necessary or requested information is displayed at any given time’ (Lidwell et al. 2003, p. 154). In the context of searching for a suitable learning path Turbeks worked out example could look like this (see Fig. 9.4): the user enters a query for a learning path aimed at attaining a certain competence. Following Turbeks example here this is done by entering keywords, but might as well take other forms like selecting from a list or pointing at an item on a competence map. The learner receives an overview of all available learning paths with the opportunity to further narrow the search by three or four parameters: e.g. study load, costs and delivery mode. Besides there is a link ‘more’ that enables the learner to set other parameters as well.

Apart from the question which criteria to include, we have to consider possible values related to these criteria. This is not an easy task as the ‘study type’ criterion in the example of Fig. 9.3 illustrates. The possible values are not self evident (cus-

Search for:		advise interior design		
<hr/>				
Focus search by:				
	<u>Study mode</u>	<u>Price</u>	<u>Recognition</u>	<u>more</u>
<hr/>				
Results	Study mode	Price	Recognition	
xxxxxxxxxx	face-to-face	320 £	yes	
xxxxxxxxxx	distance learning	400 £	yes	
xxxxxxxxxx	mixed mode	260 £	no	
xxxxxxxxxx	face-to-face	300 £	no	
xxxxxxxxxx				

Fig. 9.4 Progressive disclosure

tomised, short?) and seem in fact to represent a mixture of criteria (schedule information, workload, and study type e.g. self study). It goes without saying that values should be clear, exhaustive and mutually exclusive. Using the IEEE LOM standard (IEEE/LOM 2002) to describe metadata the learning path specification offers a set of metadata, which can be adapted to local applications without losing the advantage of interoperability by using only a subset or by defining extensions (Duval et al. 2002). This way it is possible for instance to add domain-specific information like the national legislation or health system a learning activity pertains to. Of course the benefits of using the specification are partly dependent of the scale on which the specification is adopted; an issue that will be further addressed in Sect. 9.6.

9.4.2 Following a Learning Path

According to the learning path depicted in Fig. 9.2 our learner has three options to start: the learning activity for *colour*, the learning activity for *materials*, or the first learning activity of *visual communication*. These options have to be communicated to the learner in a user-friendly way. Figure 9.5 offers an example visualisation of sequences and selections as specified by the learning path depicted in Fig. 9.2.

If the learning path of Fig. 9.2 is presented to a learner who already has mastered the competence related to *materials*, the entire path will be presented with an indication of the parts (competences) that already have been covered. In Fig. 9.5 a checked checkbox is used to illustrate this: the learner already has attained this competence. Leaving the entire learning path intact (as opposed to presenting only those competences that still have to be acquired) will support learners in developing the aforementioned structural representation of the knowledge domain. Gradually more and more activities will be marked as being completed as the learner proceeds

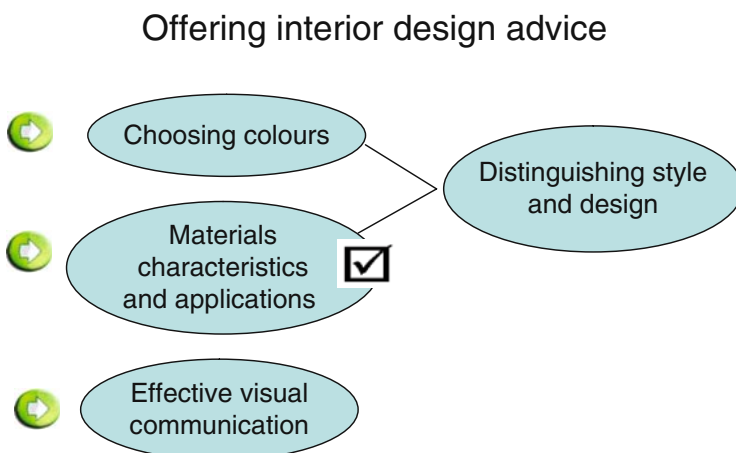


Fig. 9.5 Visualisation of sequences and selections

along the learning path. This could take place either automatically for example when the learner has passed an assessment or manually by the learner when she decides to change the activity status into completed.

So far we have discussed a single learning path, leading to the attainment of the competence 'Offering interior design advice'. Note that although it is a single learning path, there are several ways to proceed along this path: in this particular case the learner can choose a number of different activities to start with. This is what the learning path specification does: it expresses in a formal way the constraints and range of choices related to a learning path.

The actual route a learner has followed through this path is called the learning track. The learning track describes the learning activities a learner has successively completed in order to achieve the competence or learning goal associated with the learning path. A relatively simple path, consisting of a selection of two learning activities A and B, which both need to be completed, can result in two different learning tracks: [A, B] and [B, A]. In other words, a learning track describes one possible route through the learning path. The track represents a particular instance (instantiation) of the learning path. This relation between a learning path and a learning track is relevant in explaining how the learning path specification can be used to create new learning paths by learners who design their own professional competence development plan. Since the specification is a generic description of a set of learning activities, it can also be used as a template for learners to 'fill' or describe their learning track in hindsight, as the following section further illustrates.

9.5 How to Create Your Own Learning Path

Ideally there will be at least one learning path for each competence identified in a domain. However this might not always be the case: new competences will evolve over time (due to for instance technological innovations) or the available paths do not suit the learner. In absence of a suitable professional competence development programme the learning path specification can be used as a template in adapting an existing programme or creating a new one. Adapting an existing path is the simpler case: in the scenario presented above the learner might want to replace the learning activity that will lead her to acquire the sub-competence 'Distinguishing style and design by another activity with an equal result'. Instead of analysing the differences between style X and Y she would like to delve into a comparison of X and Z. In the case of a formally recognised learning path she might have to negotiate this change with a tutor. But even in the case of informal learning where the learner decides autonomously, it will be useful to actually change the learning activity description so that an accurate description of what the learner did to attain the competence can then be used for instance to store in the e-portfolio (Berlanga et al. 2008), or to present as another possible route through the learning path to learners who want to acquire the same competence (Vogten et al. 2008).

Creating an entirely new path requires for the learner to first identify the competence she wants to attain from a list or competence map or even define a new competence. If the learner selects a competence from an existing competence map, the map will offer further guidance as to identification of sub competences and relations between competences. Defining a new competence will leave things even more open, but in both cases a learning path can be described in terms of learning activities leading to the attainment of that competence. The learning path specification can be used as a template to help the learner structure and describe the learning path she intends to follow. This can be done incrementally, but could also be done in hindsight, after a learner (informally) has performed a number of activities in pursuit of developing a competence. Note that in either case – incremental and retrospective description – this entails that an actual learning track is described which can subsequently be presented as a learning path for other learners with the same goal. To this end the description of the learning track might be further enriched by identifying fixed order activities (sequences) and alternative routes (selections), but it could also be simply presented as a fixed sequence of learning activities. In the latter case the learning path and learning track completely overlap because the learning path describes only a single route. Describing informal learning paths in hindsight is more likely to provide the added value of reflection upon the learning path: some learning activities that have been done with certain objectives in mind may have proved not to be helpful at all. Of course it's highly recommended that these activities are removed from an incrementally developed description as well.

9.6 Conclusion

Describing the learning path in a formal interoperable way has several advantages even in the case of informal learning: the learning path can be stored and used in different environments for different purposes (e-portfolio, search tools, planning tools). This chapter has focused on the purposes of finding and following learning paths. In addition it has described how, in absence of suitable existing learning path, the specification can be used as a template for the adaptation of an existing learning path or the creation of a new one.

As to finding a suitable, personalised learning path a number of successive steps have to be taken:

1. Identify a learning goal. This might require some help in 'translating' a generally stated goal to a specific learning goal (competence).
2. Identify competences which the learner has already acquired.
3. Identify the available learning paths leading to the goal competence.
4. In case the number of available learning paths connected to the learning goal exceeds a number the learner can handle, a process of screening takes place in

order to reduce the set of learning paths to choose from to an amenable number. This involves specifying additional criteria for the learning path.

5. Compare the (sub)set of suitable learning paths to arrive at a choice

The learning path specification can support steps 3–5, and further supports representation of the structure and navigation options of the learning path that has been chosen. It could be argued that a suitable learning path can also be provided through recommender systems and social collaborative filtering techniques. However valuable these techniques might be to support learners' navigation through a learning path, we consider it essential that in selecting a particular learning path, learners have some notion of available options and variability. Even though navigational support is meant to help reduce disorientation as was argued in the introduction of this chapter, we still want to enable self-directed learning: 'a form of study in which learners have the primary responsibility for planning, carrying out and evaluating their own learning experiences' (Merriam and Caffarella 1991, p. 41).

An important practical implication of using the learning path specification within an educational organisation is that all learning paths – their structure and characteristics – have to be described in terms of the specification. While the advantage of metadata over free text search have been stressed repeatedly, it is also clear that there are concerns regarding the effort it takes: the accuracy required and the costs involved (Hirvasniemi and Öörni 2001; Lamminaho 2000; Pöyry Peltto-Aho and Puustjärvi 2002). In the case of the learning path specification however the costs involved in describing the metadata can be low, for several reasons. Firstly the set of metadata is relatively small. Moreover, the metadata needed are likely to exist already in the form of course catalogue information and curriculum descriptions. So the main question is how to organize this information in a way that it is stored and updated once to be subsequently used in different contexts (marketing, course catalogues, search engines etc.). Projects like the eXchange Course Related Information (XCRI 2006) project in the United Kingdom are currently working on this principle. Finally, possibilities for automated metadata generation are being investigated and some promising results have been achieved (Cardinaels et al. 2005). Automated metadata generation could be equally applied in the description of formal and informal learning paths. For an individual learner who has developed an informal learning path it will still require a considerable effort to describe it, but here too the principle of 'document once, use manyfold' holds: specifying the learning track to store it in the learners portfolio as well as submitting it to a repository open to search by others.

As was mentioned earlier in this chapter the benefits that can be achieved through use of a learning path specification depend on the scale on which the specification is adopted. For a large educational institution benefits may arrive almost immediately in the form of more automated student support. Benefits will increase considerably though in the more large scale scenario of deployment across institutions and possibly national education systems, facilitating comparison, exchange and access for (lifelong) learners. However this is not only a matter of suitable technology, but involves policy decisions on procedures for enrolment, exchange,

payments, etcetera. Both on a national and international level competence descriptions have to be compared and matched. Of course Open Educational Initiatives like OpenCourseWare (MIT 2007), OpenLearn (OUUK 2007), Curriki (2006), and many others, already indicate a shift towards opening access.

In a thought-provoking article Goldhaber (1997) pictures the shift we are experiencing from a material economy to what he calls a ‘full attention economy’. Along with this shift a breakdown of organizational barriers will take place. Others describe this shift in terms of the arrival of the ‘information age’, but as Goldhaber points out information cannot be a basis for an economy because economies are governed by what is scarce and especially these days information can hardly be considered a scarce commodity. Instead of information, attention is the scarce commodity and driving force of the new economy, where money flows to attention. Along with this shift towards a full attention economy Goldhaber envisages a breakdown of organizational barriers. The Web and other media will foster transparency and as a result, organizations will diminish in importance, relative to the importance of the individuals who are in them. And where these individuals are located will matter less and less ‘until the institution loses all coherence, all distinctness from (...) any one of hundreds of other organizations which have audiences in common’ (Goldhaber 1997).

Time will tell whether future developments will be as far stretching as this. However adoption of a learning path specification could be seen as part of the more general trend towards increased transparency.

References

- Beach, L.R.: *The Psychology of Decision Making: People in Organizations* (Newbury Park, CA: Sage 1997)
- Berlanga, A.J., Sloep, P.B., Brouns, F., Bitter, M., Koper, R.: Towards a TENCompetence ePortfolio. *Journal of Emerging Technologies in Learning (iJET)* **3**, 24–28 (2008)
- Brooks, C., McCalla, G., Winter, M.: Flexible learning object metadata. *12th International Conference on Artificial Intelligence in Education*, pp. 18–22, July, 2005, Amsterdam, The Netherlands
- Cardinaels, K., Meire, M., Duval, E.: Automating metadata generation: the simple indexing interface. *14th International Conference on the World Wide Web*, May 10–14, 2005, Chiba, Japan
- Chen, S.Y., Fan, J.-P., Macredie, R.D.: Navigation in hypermedia learning systems: experts vs. novices. *Computers in Human Behavior* **22**, 251–266 (2006)
- Curriki: Curriki Global Education Learning Community [Retrieved October 5, 2008 from <http://www.curriki.org/xwiki/bin/view/Main/WebHome>] (2006)
- Darken, R.P., Silbert, J.L.: *A Toolset for Navigation in Virtual Environments*. ACM User Interface Software & Technology (1993)
- Drachsler, H., Hummel, H.G.K., Koper, R.: Personal recommender systems for learners in lifelong learning: requirements, techniques and model. *International Journal of Learning Technology* **3**(4), 404–423 (2008)
- Duval, E., Hodgins, W., Sutton, S., Weibel, S.L.: Metadata Principles and Practicalities. *D-Lib Magazine* **8** (April 2002)
- Goldhaber, M.H.: The Attention Economy and the Net. *First Monday* **2**(4) (1997)
- Hirvasniemi, S., Öörni, K.: Educational information in the web: discussing the metadata requirements for a web service guiding citizens’ education. *International Conference on Dublin Core and Metadata Applications*, 2001, Tokyo, Japan

- IEEE/LOM: Standard for Learning Object Metadata. Learning Technologies Standards Committee of the IEEE 148.41.21 [Retrieved October 5, 2008 from http://ltsc.ieee.org/wg12/files/LOM_1484_12_1_v1_Final_Draft.pdf] (2002)
- IMS-LD: IMS Learning Design Information Model. Version 1.0 Final Specification [Retrieved February 24, 2004, from <http://www.imsglobal.org/learningdesign/index.cfm>] (2003)
- Janssen, J., Berlanga, A., Vogten, H., Koper, R.: Towards a learning path specification. *International Journal of Continuing Engineering Education and Lifelong Learning* **18**(1), 77–97 (2008)
- Janssen, J., Tattersall, C., Waterink, W., Van den Berg, B., Van Es, R., Bolman, C., et al.: Self-organising navigational support in lifelong learning: how predecessors can lead the way. *Computers & Education* **49**(3), 781–793 (2007)
- Kickmeier-Rust, M.D., Albert, D., Steiner, C.: Lifelong competence development: on the advantages of formal competence-performance modeling. *International Workshop in Learning Networks for Lifelong Competence Development, TENCompetence Conference*, 2006, Sofia, Bulgaria
- Koper, R.: Use of the semantic web to solve some basic problems in education: increase flexible, distributed lifelong learning, decrease teachers' workload. *Journal of Interactive Media in Education* **6** (2004)
- Lamminaho, V.: Metadata Specification: Forms, Menus for Description of Courses and All Other Objects. CUBER project: Deliverable D3.1 (2000)
- Lea, S.J., Stephenson, D., Troy, J.: Higher education students' attitudes to student-centred learning: beyond 'educational bulimia'? *Studies in Higher Education* **28**(3), 321–334 (2003)
- Lidwell, W., Holden, K., Butler, J.: *Universal Principles of Design* (Gloucester, MA: Rockport Publishers, Inc. 2003)
- Lueg, C., Bidwell, N.: Towards egocentric way-finding appliances supporting navigation in unfamiliar terrain. *Issues in Informing Science & Information Technology* **1**, 1097–1104 (2004)
- Martinez, P., Munday, F.: 9,000 Voices: student persistence and drop-out in further education (FEDA-report, No. 7, Vol. 2, 1998)
- Merriam, S.B., Caffarella, R.S.: *Learning in Adulthood: A Comprehensive Guide* (San Francisco, CA: Jossey-Bass 1991)
- MIT: Massachusetts Institute of Technology OpenCourseWare [Retrieved October 5, 2008 from <http://web.mit.edu/ocw/>] (2007)
- OUUK: Open University of the United Kingdom OpenLearn [Retrieved October 5, 2007 from <http://openlearn.open.ac.uk/>] (2007)
- Pöyry, P., Pelto-Aho, K., Puustjärvi, J.: The role of metadata in the CUBER system. SAICSIT 2002 South African Institute of Computer Scientists and Information Technologists on Enablement through Technology, Port Elizabeth, South Africa (2002)
- Rundle-Thiele, S., Shao, W., Lye, A.: Computer Process Tracing Method: Revealing Insights into Consumer Decision-Making. Australian and New Zealand Marketing Academy Conference (ANZMAC) [Retrieved October 5, 2008 from <http://smib.vuw.ac.nz:8081/WWW/ANZMAC2005/cd-site/pdfs/3-Consumer-Beh/3-Rundle-Thiele.pdf>] (2005)
- Schwartz, B.: *The Paradox of Choice. Why More Is Less*, 1st ed. (New York: HarperCollins Publishers 2004)
- Simpson, O.: Access, Retention and Course Choice in Online, Open and Distance Learning. Third Eden Research Workshop, Oldenburg, Germany, 2004
- Turbek, S.: Boxes and Arrows: The Design Behind the Design [Retrieved July 2, 2008 from <http://www.boxesandarrows.com/view/advancing-advanced>] (2008)
- Vogten, H., Koper, R., Martens, H., Van Bruggen, J.: Using the personal competence manager as a complementary approach to IMS learning design authoring. *Interactive Learning Environments* **16**(1), 83–100 (2008)
- XCRI: eXchanging Course-Related Information [Retrieved October 16, 2006 from <http://www.elframework.org/projects/xcri>] (2006)
- Yorke, M.: *Leaving Early. Undergraduate Non-completion in Higher Education* (London: Falmer Press 1999)

ture for interoperable e-assessment services that support these new assessment types. The work is a follow-up on the assessment model developed in a joint project of CELSTEC (formerly OTEC) and Cito (Joosten-Ten Brinke et al. 2007). This *conceptual model* was a first step towards an interoperability standard for assessment that would go beyond the QTI specification in its support for a variety of assessment forms, as well as in the definition of assessment plans that allow long term monitoring of competence growth. The work reported in Chap. 10 takes the model considerably further by embedding the elements of the conceptual model in APL, assessment process modelling language, that is activity- and process centred. The chapter demonstrates the modelling of a performance assessment and demonstrates execution of the APL in an LD and QTI-compatible run-time environment.

The second service presented in this section offers support in *placement*. Professionals that enter a learning network already acquired competencies in different settings that range from formal settings (school, university) to their professional environments. The placement service supports professionals in choosing appropriate routes in competence development by mapping existing qualifications, irrespective of their origins, onto learning activities that lead to a competence goal, in order to detect and remove redundant learning activities. In Chap. 11 placement services are described from the perspective of Accreditation of Prior Learning (APL). To support APL a gamut of services can be conceived (see Joosten-Ten Brinke et al. 2008 for examples) in various stages of an APL process. For Learning Networks a placement service is designed and developed to *assess the evidence* in the APL portfolio and to help decide whether the outcomes of learning activities have already been met. The challenge that placement services in Learning Networks have to meet is that in Learning Networks we cannot assume a joined, controlled vocabulary with which the contents and outcomes of learning activities and the contents of the portfolios of the members of a learning network are described. For that reason placement requires mechanisms independent from any type of formal data (Kalz et al. 2007). Chapters 11, 12 and 13 describe the background, tools and techniques and a validation scenario for a placement service. The latter includes a report of successes, failures and lessons learned.

Chapter 10 on assessment services is of particular interests for researchers and developers working in the areas of learning standards, in particular IMS Learning Design and QTI.

Chapters 11, 12 and 13 are discussing placement services in Learning Networks. Chapter 11 gives the theoretical background of the content-based placement service and Chapter 13 details the work in validating such a service in terms of the quality of its analysis and recommendations. These chapters are of particular interest to researchers and to those who are dealing with implementation and evaluation of this types of services. Chapter 12 provides an overview of tools and techniques and is meant for designers and developers in particular.


References

- Joosten-ten Brinke, D. et al.: Modeling assessment for re-use of traditional and new types of assessment. *Comput. Hum. Behav.* **23**, 2721–2741 (2007)
- Joosten-Ten Brinke, D., et al.: The quality of procedures to assess and credit prior learning: implications for design, *Educ. Res. Rev.* **1**, 51–65 (2008)
- Kalz, M. et al: Positioning of learners in learning networks with content, metadata and ontologies. *Interact. Learn. Environ.* **15**, 191–200 (2007).
- Koper, E. J. R., Rusman, E., Sloep, P. Effective learning networks. *Lifelong Learn. Eur.* **1**, 18–27 (2005)

Chapter 10

A Process-Oriented Approach to Support Multi-Role Multi-Stage E-Assessment: A Case Study

Yongwu Miao and Henry Hermans



10.1 Introduction

This chapter presents a process-oriented approach to support assessment in learning networks. Assessment in its broadest sense can be conceived as the ‘systematic gathering of information and evidence about a person’s properties (qualifications), based on a process, a product or the progress of a person, for the purposes of certification, placement or diagnoses in formative and summative contexts’ (Joosten-ten Brinke et al. 2007). This definition includes classical assessment and complex assessment. Examples of classical assessment are multiple-choice, fill-in-the-blank, open question, matching, and ordering. The information gathered in classical assessment is usually quantitative, numerical data, which can be judged against a simple, pre-set standard. Examples of complex assessments are peer assessment, 360 degree feedback, performance assessment, and portfolio assessment. A complex assessment usually involves multiple people with diverse roles who perform various activities in the same stage or different stages in sequence and/or in parallel in a coordinated way. By addressing complex learner traits, complex assessment aims to foster

Y. Miao (✉)

Centre for Learning Sciences and Technologies, Open University of the Netherlands,
Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands
e-mail: yongwu.miao@ou.nl

deep learning and the development of competences (Boud et al. 1999; Gipps 1999; Topping 1998).

In order to support assessment in learning networks, it is required that software tools and services for conducting e-assessment are being provided and many software tools have indeed been developed to support e-assessment. We can classify these software tools along two dimensions: (1) their support of classical or complex e-assessment and (2) their support of interoperability and reusability.

Most assessment software tools, such as Zoho Challenge fall into the first quarter. The tools in this category focus on individual assessment, in which the computer takes over the role of providing feedback and evaluation. They support authoring, answering, and evaluating classical questions and they have their own data structures. That is, the defined questions and assessment results can not be shared by other assessment software tools. This alone makes them not suitable for use in learning networks. The software tools in the second quarter can support in particular peer assessment, for example SPARK (Freeman and McKenzie 2002), Peer Grader (Gehring 2001), and Espace (Volder et al. 2007), and 360 degree feedback as eSPRAT (Lockyer 2003) does. However, they have the same limitations in interoperability and reusability.

IMS Question and Test Interoperability (QTI 2003), an international e-learning technical standard has been designed with the special purpose to support interoperability of assessment tools and reusability of e-assessment resources. IMS QTI describes a data model for the representation of assessment items and tests and the reports of their corresponding results. Almost all classical question types can be supported by QTI. Moreover, it provides sufficient flexibility for further development towards the advanced constructed-response items and interactive tasks that is envisaged as the future of assessment (Almond et al. 2001). Examples of QTI-compatible software tools are APIS (APIS 2008), AQuRate (AQuRate 2008), and R2Q2 (R2Q2) (Wills et al. 2006).

However, QTI provides insufficient expressiveness for modelling a complete process of complex e-assessment that goes beyond the selection of certain elements from a pre-defined set and/or assigning or constructing a simple entity or relation. QTI ignores who will be involved in an assessment and what roles they play; what kinds of activities should be performed by whom and in which sequence; what assessment resources will be produced and used in the assessment process, and what dynamic changes in the process may happen and under which conditions. In short,

Table 10.1 A classification of assessment software tools

Classical e-assessment	Complex e-assessment	
I Zoho Challenge	II SPARK, Peer Grader, Espace and eSPRAT	Can not support interoperability and reusability
III APIS, AQuRate and R2Q2	IV None	Can support interoperability and reusability

it provides insufficient support for representing a *multi-stage and multi-role assessment*. As illustrated in Table 10.1, there is no tool in the fourth quarter supporting interoperability and reusability of tools and resources of complex e-assessment.

We aim to provide technical support for complex e-assessment in learning networks in an interoperable and reusable manner. We build on our experience with the development of the IMS Learning Design specification (IMS LD 2003), a standard educational modelling language used to specify a wide range of pedagogical approaches and strategies (Koper and Olivier 2004). From here we set out to develop an assessment process modelling language that can be used to, at least, specify a wide range of assessment approaches and strategies (including classical assessment). The starting point of the development of the assessment process modelling language is the OUNL/CITO model, a meta-model created and validated in a joint effort by the Open University of the Netherlands (OUNL) and the Citogroep (CITO) (Hermans et al. 2006; Joosten-ten Brinke et al. 2007). This meta-model provides the ‘ingredients’ to prepare a variety of assessments ‘recipes’. It has been set up to model classic, as well as complex assessment ‘recipes’ like performance assessment, peer assessment, or 360 degree feedback. However, the OUNL/CITO meta-model provides only a *conceptual framework* for modelling and analyzing assessment from a perspective of assessment experts. An assessment model represented in the OUNL/CITO meta-model lacks sufficient operational semantics and cannot be executed by a workflow engine. For automating a complex assessment process in learning networks, an assessment process model has to be formally specified by explicitly modelling activities performed by multiple roles and people, their temporal relationships, and the transfer of information between activities. Therefore, we developed an assessment process modelling language by incorporating concepts inspired by the OUNL/CITO model into the framework of IMS LD (Miao et al. 2008a). The assessment process modelling language defines a set of assessment-specific language constructs, which can be used to describe an assessment approach or strategy as a formal assessment process model. Finally, a generic execution environment, which is compliant with the assessment process modelling language, should be available to support the execution of the formal assessment process models.

Figure 10.1 provides an overview of a process-oriented approach to support complex e-assessment. By adopting this approach, the assessment designer (planner or organizer) can consider the design issues based on the OUNO/CITO meta-model when she or he designs an assessment plan at design-time. The assessment plan can be described in terms of OUNL/CITO meta-model. Then the author (note that sometimes the designer and the author are the same person or the same group) can develop a computerized assessment plan as a formal assessment process model by using the assessment process modelling language. Then, the formal assessment process model will be used in practice as a guideline to organize and structure how activities should be performed, by whom, and in what order. In particular, a language-compatible execution environment can interpret the formal assessment process model and orchestrate activities that are performed by different participants in different stages at run-time. In addition, the same formal assessment process

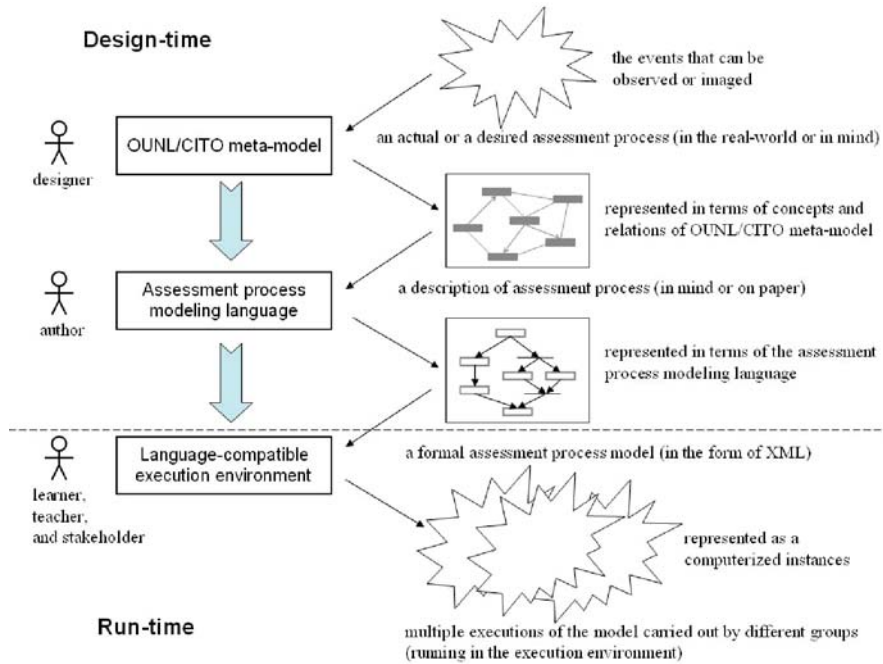


Fig. 10.1 An overview of the process-oriented approach to support e-assessment

model can be instantiated multiple times to support different groups and it can be modified to support a customized assessment.

The benefits of this approach are:

1. The OUNL/CITO meta-model can help designers to systematically express general design ideas about the assessment to be modelled.
2. The assessment process modelling language can facilitate representation, understanding, communication, comparison, and reuse of a variety of assessment practices in a formal way.
3. An execution environment can guide and scaffold learner, tutors, and other stakeholders to perform activities by providing guidance and awareness information, by configuring a workspace for carrying out prescriptive tasks, by controlling and changing the sequence of activities based on the execution state and circumstantial information, and by orchestrating the efforts made by different roles/participants.
4. It will not be necessary to develop dedicated software tools for particular assessment approaches. A generic execution environment can be used for the execution of a variety of assessment types and strategies.

In this chapter, we first describe an authentic exam. Throughout this chapter this exam is used as an exemplary assessment process. We will use it to illustrate how to model an assessment process based on the OUNL/CITO meta-model, how to author a formal assessment process model using the assessment process modelling language, and how to execute a formal assessment process model in a run-time environment.

10.2 A Use Case of Performance Assessment

In this section we present an authentic exam to illustrate our approach. The following description is an accurate account of an exam that took place at a school of arts. Further on in this chapter this use case will be extended to the requirements for computer supported executing in learning networks.

After four years of studying, Maaïke is ready to take her second grade violin exam. This exam consists of playing three selected violin parts, accompanied by the piano. These three parts will be judged (rated) by an examining board of three persons, led by a chair. The Chair's weight in the rating process is higher than that of the other board members.

The process of the exam is described in the narrative below:

1. Maaïke plays three assigned violin parts and the performances are observed by three raters.
2. Each performance will be judged by three raters using an assessment form.
3. The chair collects all rubric values provided by all raters and calculates the weighted rubric values/scores, weighted indicator scores, and qualification scores. In the end, a decision will be made according to decision rules.
4. Maaïke is informed about the decision: 'grant certificate' or 'schedule for another exam'.

When a rater judges Maaïke's performance, an assessment form (see Table 10.2) is used. Each rater will assign rubric values in the cells marked with the ovals. All ratings are to be given on a scale from 1 to 10. Certain qualification values (e.g., 'Instrumental-auditive use') will be calculated as an average means of relevant rubric values (e.g., 'Lead-play and repeat by ear' and 'Improvisation', in this case) as indicated by the arrows.

To pass the exam and obtain a certificate, all six categories (qualifications) must have a mean score of at least 6. A score of 5 on one category is allowed, provided this is compensated with a higher score on one of the other categories. Besides the scores and decision, final remarks may be added to the certificate.

For a better understanding, we depict the exam process as a semi-formal model in a UML activity diagram. As shown in Fig. 10.2, a swimming lane shows all activities (represented in the round-corner rectangle) performed by a participant. The solid arrow indicates the control-flow among the activities. Some activities are

Table 10.2 Assessment form used in the example

	Rating	Total
Instrumental-auditive use	-	-
Lead-play and repeat by ear	-	-
Improvisation	-	-
Tonal mastery	-	-
Quality of tone and sound	-	-
Tonal purity	-	-
Technical skills	-	-
General	-	-
Scales and chords	-	-
Rhythmical skills	-	-
Ensemble	-	-
Expression skills/presentation	-	-

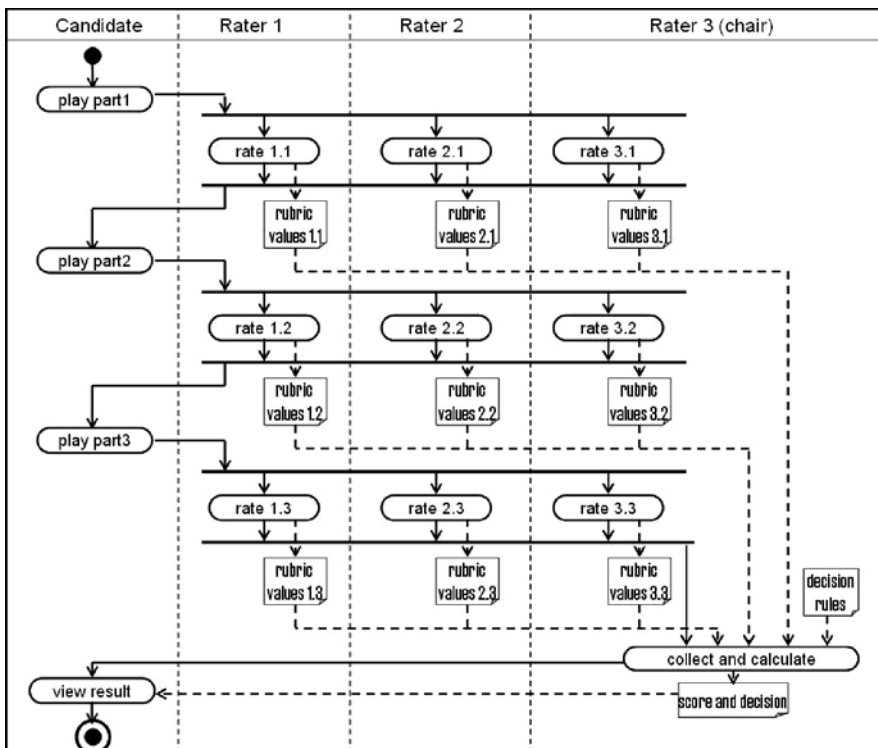


Fig. 10.2 The semi-formal assessment process model of the exam in a UML activity diagram

performed in parallel. Certain information such as rubric values and decision are represented as notes and the dashed arrows indicate the information-flow from/to activities.

From this exam we can analyze characteristics of complex assessment. In contrast to classical assessment, where there is usually one rater only (a teacher or even a computer), there may be several types of roles (e.g., peer, tutor, and even customer) and numbers of raters involved. Raters may assess different aspects of a candidate's performance as in 360-degree assessment. In contrast to classical assessment, where the candidate usually answers a list of standard types of questions, the candidate may need to demonstrate competences by performing various tasks using assessment-specific tools (e.g., concept-mapping tool and simulators) or even using physical tools in real-world. In contrast to classical assessment, where quantitative data usually serve as evidence and standard/correct answers are available, the information provided by the candidate may be tangible (e.g., a concept-map or a design) or intangible (e.g., performance of music/dance or playing a car-driving simulator). The evidence gathered for assessment may be qualitative, descriptive, and narrative information. There is no standard/correct answer and the evidence has to be interpreted by raters.

10.3 Modelling Assessment Based on OUNL/CITO Model

In this section we briefly present the OUNL/CITO model and illustrate the model through the use case in Sect. 10.2. The OUNL/CITO model is an extensible educational meta-model for assessment. The meta-model allows a tight embedding of assessments in educational practice and it caters for new types of observation and interpretation. The model is built on several sub-models, each one matching a different stage in the assessment process. Within the assessment process six stages can be distinguished: (1) assessment design, (2) item construction, (3) assessment construction, (4) assessment run, (5) response processing, and (6) decision making. In this section, we present the main concepts of the OUNL/CITO model according to the six stages. A detailed description of the model can be found in Joosten-ten Brinke et al. (2007).

10.3.1 Assessment Design

The reasons for using assessments are expressed in the stage of assessment design. The challenge in assessment design is to select the assessment types that yield the appropriate evidence of the learners' competence, skills or knowledge. The result of work in the stage of assessment design is an object called the *AssessmentPlan*, which describes the:

- *function* of the assessment (diagnosis, placement, certification);
- *characteristics* of the target group;
- *qualifications*, as well as qualification levels;
- *indicators* for the qualifications and their weights;
- *decision* rules to be used;
- (possible)*scenario(s)* for navigating through the assessments and gathering the evidence to prove the mastery of the Qualification(s) addressed.

In the use case the AssessmentPlan shows the following characteristics:
 the assessment function is certification;
 the target group are violin students studying at the school of music, and who have passed their first grade violin exams;
 the qualifications are: instrumental-auditive use, tonal mastery, technical skills, rhythmical skills, ensemble, expression skills/presentation.

QualificationLevel: 2

the decision rules: see the use case: there is only one UnitOfAssessment, covering all Qualifications (some of them based on several Indicators) to be rated and scored.

10.3.2 Item Construction

In this stage the *qualifications* and their *indicators* give direction to the construction of Items. Often these qualifications may not be observed directly, and indicators have to be specified: these are measurable elaborations of the qualification. A score on an assessment has a meaning for a qualification, but it is directly based on scores on the underlying indicators by applying a calculation rule on the scores. Note that the term *item* in this model has to be interpreted in a broad sense. For example, it applies to a multiple-choice item with four answering options, as well as to a task in which a candidate has to show a performance.

An *item* is defined as the smallest unit to estimate one or more *qualifications* of a *person*. An Item may address one or more *indicators*.

An *item* consists at least of the following four components:

- *Prompt*: the stimulus, task, or activity that evokes some behaviour of the Person.
- *ResponseMode*: description of the way the Person is intended to respond to the Prompt.
- *Description of TargetGroup*.
- *RatingPrescription*, to be specified for all Rater roles and Indicators. These RatingPrescriptions in turn consist of one or more rating rubrics (Rubrics); for each Rubric a scale has to be specified and rules must be present on how to transform

an ItemResponse into a specific scale value. Furthermore rules must be present for the transformation of these scale values into scores.

Optional parts of an item are:

- *Feedback*: recommendation or information provided to the Person based on gathered ItemResponse(s);
- *Hint*: a specific item instruction, e.g. partial clue, available to the Person to make it easier to give the right response;
- *Context*: a piece of discourse, material or case description to which items are connected within an assessment.

In the use case each part Maaike has to play is considered an Item. Each Item is used to estimate the whole range of Qualifications from the AssessmentPlan. The Prompt for each Item is provided by the chair; the ResponseMode is 'Demonstration'. The RatingPrescription in this example is not present as a physical object, but is implicitly present in the minds of the Raters. There are two Raters roles: 'member' and 'chair'.

10.3.3 Assessment Construction

The central concept within this stage is the UnitOfAssessment. A UnitOfAssessment is derived from a UnitOfAssessmentDefinition (part of the AssessmentPlan) and contains the collection of Items to be presented to the Person. A UnitOfAssessmentDefinition defines several rules. These rules are about the composition of the assessment, rules prescribing what items may be used and in what order and rules that specify how the final score on a unit of assessment will be calculated. The definition defines which trait will be assessed in a specific unit of assessment (unit of assessment trait) and which indicators are used to this purpose (unit of assessment indicator).

Based on one UnitOfAssessmentDefinition ('blueprint' of UnitOfAssessment) several different UnitsOfAssessment may be generated (e.g. using an item bank system) or developed (e.g. by a teacher or expert), each of them having the same objectives and the same Indicators, but possibly containing different Items (AssessmentItems).

In the use case of the grade violin exam the UnitOfAssessment consists of three Items (the three violin parts to be played). These Items have been selected by the violin teacher (who happens to be one of the Raters).

10.3.4 Assessment Run

After a `UnitOfAssessment` has been created (based on the `UnitOfAssessmentDefinition`), it can be run or executed for an individual `Person` or group of `Persons`. The data regarding all `Persons` taking part in a specific assessment session are recorded in the `AssessmentSession`. Specific data about single `Persons` (like role or presentation medium) are recorded within the `AssessmentTake`.

An `Item` prescribes how a `Person` is expected to respond: picking an answer from a predefined set of options (selection), actively producing a response (construction) or demonstrating a task, procedure or activity while being observed (demonstration). The class `ItemResponse` records this response and always addresses one `Person` or one group of `Persons` and one `Item`.

In the use case the `UnitOfAssessment` consists of three `Items` selected by the violin teacher. For the `Person` there are no specific `presentationMedia` of roles to choose from. Characteristic for the `ItemResponses` is that they are transient. They haven't been archived or recorded

10.3.5 Response Processing

The stage of `ResponseProcessing` can be split up into two separate sub processes: rating and score computing. The rating process is concerned with assigning one or more rates to an `ItemResponse` of a `Person`, and may be executed by different types of `Raters` (like teacher or peer), using the `RatingPrescription(s)` specified. Furthermore a role (like 'peer') may be fulfilled by a variety of persons. The `RatingPrescription` allows for assigning a lower and upper limit to the number of raters in a specific role (`minNumber` and `maxNumber`).

The provided ratings for `ItemResponses` are stored in `RubricValues` and transformed into scores (`RubricScores`), according to the `RubricScoringRules` as specified for the `Item`.

When for a `Person` the `RubricScores` are complete within a `UnitOfAssessment`, these scores can be transformed into `IndicatorScores` for the `Indicators` present within the `UnitOfAssessment`. In this transformation the weight of the `Rater` types is to be taken into account.

When all `IndicatorScores` within a `UnitOfAssessment` are available, the `QualificationScores` for this `UnitOfAssessment` can be computed, of course taking the weight of the `Indicators` specified for the `Qualification` into account.

When all `IndicatorScores` within an `AssessmentPlan` are known for a `Person`, the `QualificationScores` can be computed. These scores serve as input for the `Decision-Making` process.

The processes of taking the assessment and rating run almost parallel in this example. The Raters are physically present at the exam where Maaïke has to demonstrate her skills. There is no physical registration (audio, video) of the ItemResponses to be rated afterwards. This creates additional conditions for the rating process.

The rating in this use case takes place in the absence of explicit RatingPrescriptions with rubrics, matching scales, scale values and scoring rules.

This does not imply that the Raters do not use such rubrics, but that these are implicitly present in the Raters' heads.

Furthermore there is no registration of the rating of the separate Raters of the exam committee. In close consultation two IndicatorScores are agreed upon for each of the first three Qualifications. For all these three Qualifications the Indicators have the same weight. The QualificationScore has been obtained by averaging.

There are no Indicators present for the remaining three Qualifications. How the QualificationScores have been reached remains implicit.

The IndicatorScores as well as the QualificationScores have been represented on a scale ranging from 1 (lowest) to 10 (highest).

In this example the raters' experience and reputation make up the quality of the rating.

10.3.6 Decision Making

The last phase in the assessment process is taking decisions. The type of decision, such as diagnosis, placement or certification, has been specified within the AssessmentPlan, together with the DecisionRules to be used.

In the use case the chair of the committee takes the decision whether to certify Maaïke. To obtain the certificate Maaïke must at least have a score of 6 or higher for each of the qualifications.

A score of 5 on one category is allowed, provided that it is compensated by a higher score on at least one of the other categories.

Maaïke is informed about the decision regarding her violin exam by the chair who also presents a short explanation.

10.4 Authoring an Assessment Process Model with an Assessment Process Modelling Language

In this section we briefly introduce the assessment process modelling language and then illustrate how to specify the use case using the assessment process modelling language. As stated before, the OUNL/CITO model provides a conceptual framework to view a whole process of assessment from design stage to the decision-making stage. However, it cannot be directly used to support computer-supported or blended assessment processes. Conceptually speaking, like QTI, the OUNL/CITO model is not an activity-centred and process-oriented model. Important information about control flow and information flow from one activity/role to another is not present in the OUNL/CITO model. A description of an assessment process based on the OUNL/CITO model cannot be executed by a process enactment service. In order to support automation of complex assessment in learning network, an assessment process modelling language has been developed, in which the ‘ingredients’ of the OUNL/CITO model have embedded in the framework of IMS LD, intended to be interpreted and handled by machines.

10.4.1 Assessment Process Modelling Language

Seven stages are defined in the assessment process modelling language as shown in Fig. 10.3. The stage *design* is defined for representing all design work to be done at run-time. The most important tasks done in assessment are gathering evidence and making judgment. We distinguish between these work stages using the terms *evidence creation* and *assessment*. We merge the stages of response processing and decision making by using a more generic term *processing* to specify all work done by participants or by machine in this sub-process, for instance, creating or/changing information, format or media to present the assessment result. Sometimes (e.g., in a peer assessment) the candidate has a chance to react on the feedback and to ask assessors for elaboration of comments. We define a separate stage called *reaction* for specifying this sub-process. In addition, the assessment result and decisions should be sent to the candidate or other stakeholders. The stage *information* is designed for that purpose in the modelling language. Finally, both the start point and end point of integrated learning, teaching, and assessment are often a learning or teaching stage. In summary, seven types of stages constitute an assessment process model. As shown in Fig. 10.3, they are: learning or teaching, design, evidence creation, assessment, reflection, process, and information.

They may be executed in different sequences in different models and not all stages need to be included in an assessment process model. In a particular case, a teacher may grade students based on memory and then an evidence creation stage can be excluded. In contrast, a stage may be repeated one or more times. For example, further evidence may need to be gathered after an initial assessment; and even an additional design stage may be needed for creating additional assessment items

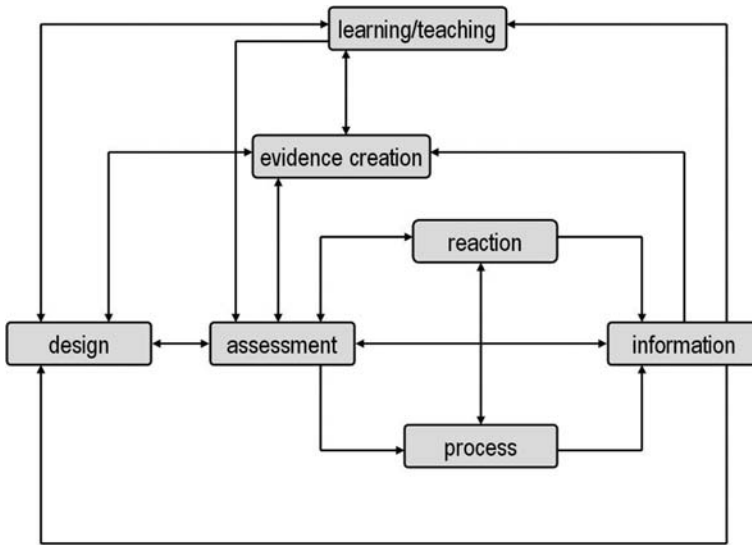


Fig. 10.3 Process structure model

according to the student’s response at run-time. Sometimes a peer assessment can be designed in a way that enables the assessee to review the feedback and request for elaboration. The assessor may provide further comments and detailed explanations. In some complicated cases, multiple loops may be defined in an assessment process model. Therefore, many concrete assessment process models can be derived from this generic process structure model.

In each particular type of stage, only certain types of activities can be performed. An *activity* is a logical unit of a task performed individually or collaboratively within a stage. Figure 10.4 shows which types of activities can be performed within a certain stage. For example, *designing demonstration assignment*, *constructing assessment items/test*, *design assessment form*, and *design rules* can only be specified in the *design* stage. In the *evidence creation* stage, *editing portfolio (items)*, *editing evidence*, *responding to assessment test/item*, *demonstrating*, and *interview* are allowed. In addition, more than one activity can be performed within the same stage. A set of activities can be grouped as an activity-structure. Four types of activity-structures are specified: sequence-structure (all activities will be performed in a prescribed sequence), selection-structure (a given number of activities selected from a set of candidate activities will be performed in any order), concurrent-structure (a set of activities are performed concurrently by the same individual or different users), and alternative-structure (one or more activities selected from a set of candidate activities according to a prescribed conditional expression will be performed). A stage, an activity-structure, and an activity have common attributes such as *completion-condition* (e.g., user-choice, time-over, artefact-available, and even user-defined conditions) and *post-completion-actions* (e.g., show/hide information/activity).

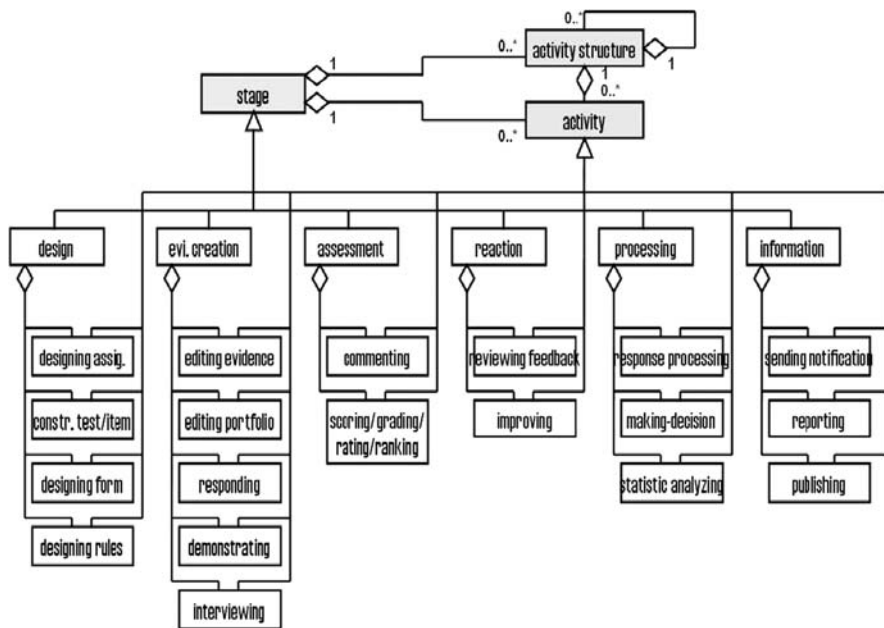


Fig. 10.4 Relations between stages and activities

An activity is performed by people with a certain role. *Role* is used to distinguish different types of participants in an assessment process. A user is allowed to have several roles at the same time and many users can play the same role in an assessment process. A role has to perform certain types of activities within an associated stage. For example, an *assessee* can perform a *responding* activity within an *evidence creation* stage and an *assessor* can perform a *commenting* activity within an *assessment* stage.

An activity may have input and output. *Artefact* is used to represent a tangible object (e.g., a paper-based questionnaire and a digitalized concept-map) or an intangible object (e.g., conceptual thinking and observed performance/behaviour) created, introduced, and shared within and/or across activities as an intermediate product and/or a final outcome. *Information resource* is another kind of input of an activity. It differs from an artefact because it is available and remains unchanged during the entire assessment process. For example, in an *evidence creation* stage, an *assessee* is required to write an analysis report about an article, that will be available in the assessment environment and is left unchanged in the assessment.

People usually need a certain service to produce or consume an artefact when performing an activity. *Service facility* is used to specify the type of computer-supported service (e.g., QTI tool, concept-mapping tool, and chat) or a human-provided service (e.g., piano companion) for handling certain types of artefacts (e.g., QTI items, concept-map, and music) or/and for facilitating communication and collaboration.

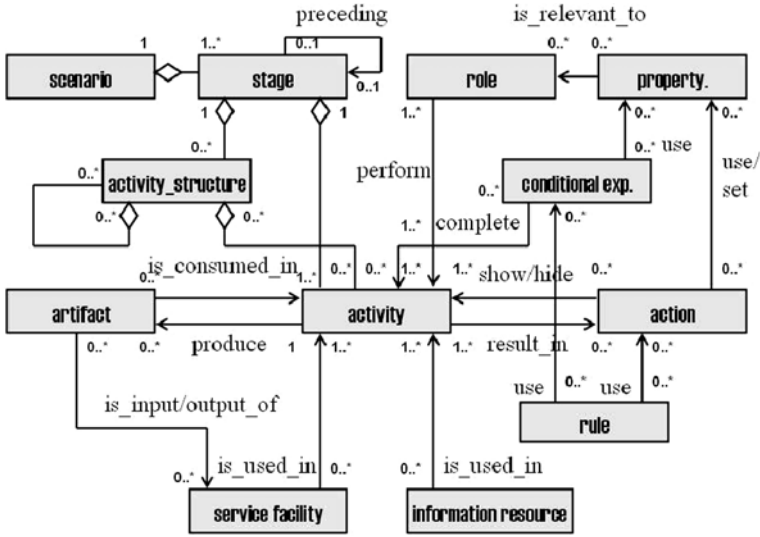


Fig. 10.5 Conceptual structure model

Property is designed to capture any information relevant to the process or to certain roles. Examples of properties are a process status and a final score.

Rule consists of conditional expressions and a set of actions and/or embedded rules in a form: If (conditional expression) Then (actions) Else (actions/rules). A *conditional expression* is a logical expression on the attributes (e.g., assessment-type, activity-status, user-in-role, role-in-activity, artefact-default-service, and etc.), artefact, and properties. An *action* is an operation performed by the system. Exemplar actions are *change attribute values* and *show/hide activity/information-resource*. Thus, a rule can be used to model dynamic features and support adaptive assessment.

Figure 10.5 illustrates the main structural relations between the concepts described above. In general, an assessment process model specifies: following certain *rules* people with various *roles* perform assessment-relevant *activities/activity-structures* allocated to them; they do so in associated *stages* using *service facilities* and *information resources* in order to consume and produce *artefacts*. To define an activity, we need to specify the roles involved, input and output artefacts, services needed, information resources referred to, completion-conditions, post-completion-actions, etc.

Figure 10.6 shows the levels of semantic aggregation in the framework of IMS LD. The semantically highest level is *assessment design*, which aggregates a collection of *components* and a *method*. A component can be one of five different types: *role*, *artefact*, *service facility*, *information resource*, and *property*. More detailed categories of each component are also depicted in Fig. 10.6. A *method* consists of one or more *assessment scenarios* and a set of *rules*. An *assessment scenario* con-

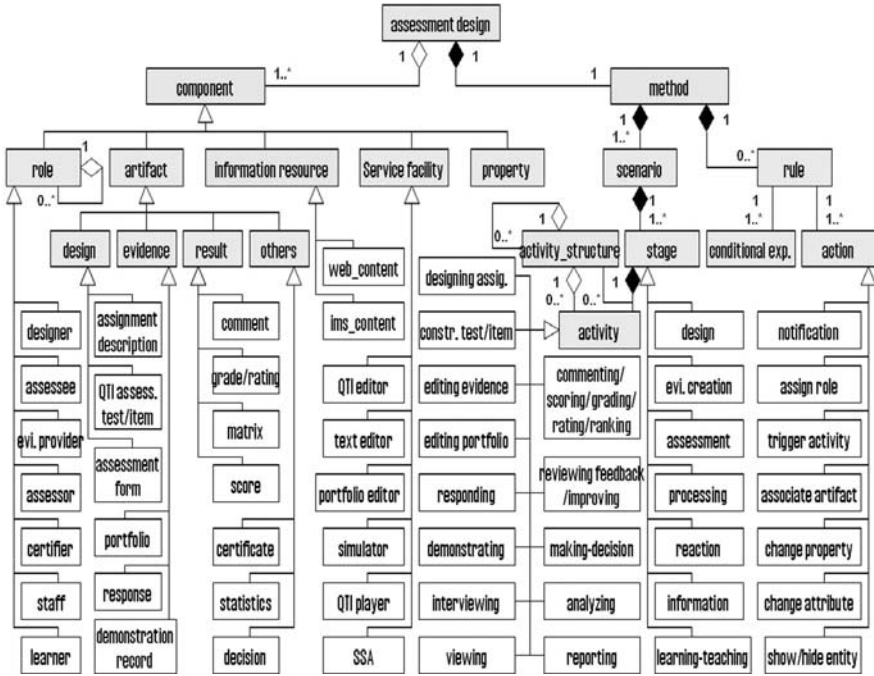


Fig. 10.6 Semantics aggregation model

sists of one or more sequential *stages*. Each *stage* consists of a set of *activities* and/or *activity-structures*. A *rule* consists of a set of *conditional expressions* and a set of *actions* in a structuralized if-then-else/else-if format. More detailed description can be found in Miao et al. (2008a).

A formal assessment process model is a package of assessment design with associated information resources and services. Assessment design has attributes such as identifier, title, description, assessment objectives, assessment types, etc. The attribute description can be used to informally present information such as assessment function, target group, estimated duration, and design rationales. *Assessment-objective* is used to describe the intended outcome of the assessment in terms of information resource or competence proficiencies. *Assessment-type* is used to define a way to yield and evaluate evidence. The possible choices are classical test, self-/peer assessment, portfolio assessment, 360 degree feedback, etc. Each choice will provide additional restrictions to the conceptual model. A detailed example of a peer assessment is given in Miao et al. (2008b).

With the help of semantically familiar vocabularies, the assessment process modelling language makes it easy and intuitive for practitioners to specify their own assessment process model.

10.4.2 *Specifying a Formal Assessment Process Model*

In this sub-section, we illustrate how to specify a formal assessment process model using the assessment process modelling language. A descriptive model can be developed by referring to Maaïke and her classmates' exam. We can create a generic assessment process model for all students in Maaïke's class. We define 'Candidate' as a kind of assessee to replace Maaïke. Such a model can be customized and reused for the exams in the future. The procedure to specify a formal assessment process model with the modelling language consists of six steps.

Specifying a process structure: There is no a standard assessment process model which can be used for all assessment. Thus, the assessment designer has to specify which stages have been (or should be) undergone and in which sequence. Because the use case is a pure assessment process which does not align with learning or teaching processes, the learning/teaching stage will be not included in this assessment process model. Because assessment design work (e.g., the three assigned violin parts, decision rules, and an assessment form) has been done before the exam starts, the assessment design stage will be not included in the model. In the exam a candidate has no chance to react on the score, review the feedback, or request for elaboration, the reaction stage will be excluded. Therefore, when we model the exam, the starting stage is one of evidence creation in which the candidate plays the first part. After that, three raters will evaluate the performance of the candidate and assign rubric values of the first part in an assessment stage. Then the candidate and raters will repeat evidence creation and assessment for two additional rounds. After all three parts are finished, the chair will collect all rubric values and calculate the final score in a processing stage. In a computer-supported assessment environment, this task can be done by the machine. Finally, the candidate will be informed about the final score and the decision in an information stage.

In summary, the assessment process model describing the exam consists of a sequence of stages: evidence creation #1, assessment #1, evidence creation #2, assessment #2, evidence creation #3, assessment #3, processing, and information.

Specifying activities for roles in each stage: We should specify who performs that activities in each stage.

1. In evidence creation 1, the candidate plays the first part while three raters observe the performance.
2. In assessment 1, three raters assign rubric values for the performance of the first part.
3. In evidence creation 2, the candidate plays the second part while three raters observe the performance.
4. In assessment 2, three raters assign rubric values for the performance of the second part.
5. In evidence creation 3, the candidate plays the third part while three raters observe the performance.

6. In assessment 3, three raters assign rubric values for the performance of the third part.
7. In processing, the machine automatically collects rubric values and calculates the qualification values and final score and makes decisions according to decision rules.
8. In information, the candidate will be informed about the final score and whether she/he can obtain a certificate.

Specifying services and recourses used to perform each activity: When playing each part, the candidate is accompanied by piano music played by a pianist. We can image that it is possible to use a DVD-based piano accompaniment as a kind of service for demonstrating the competences to play the violin. In addition, the information about which part should be played is available before playing. In each rating activity, an assessment form (see Table 10.2) should be available so that the raters can assign rubric values for the performance of each part. This assessment form can be defined using a QTI authoring tool.

Specifying process control-flow: It is necessary to specify how an activity and a stage will be finished and how to trigger the succeeding stages and activities. *In each evidence creation stage, when a candidate finishes playing, the stage will be terminated.* Other options could be ‘taking ten minutes’, or ‘decided by the chair’. Usually, when all activities within a stage are terminated, the stage will complete and the succeeding stage starts automatically.

In each assessment stage, three raters assign rubric values using a paper-based assessment form (see Table 10.1). In a computer-based assessment environment, each rater can use a desk-top or a PDA to assign the rubric values. When three raters finish assigning rubric values for each part, the stage completes. As specified, there are three loops between evidence creation and assessment in this model.

After three loops, the collecting and calculating activity will be active in the processing stage. In the original exam, the chair collected and calculated manually, but in a computer-based assessment environment this activity can be performed automatically by the system. When all indicator scores are calculated the chair can terminate the activity. After that, the final score and the decision will be created according to the rules, which will be specified below. The termination of this activity will trigger the start of the informing activity in the information stage. In the original exam, the chair informed the candidate face-to-face. In a computer-based environment, this informing activity can be carried out by sending a notification (via e-mail or handy message) or by displaying the result in a public large screen on the wall.

Specifying artefacts and artefact-flow: Here we identify and specify what is produced in each activity and whether it will be consumed in other activities. In each activity of violin playing an intangible artefact is created as performance output. In the original exam, three raters observed and used the performance as input for their rating activities. We can image that the exam is conducted in a computer-supported environment in a remote and/or asynchronous manner. A video/audio clip can be

captured as digitalized performance, which can be archived and used for assessment. The artefact produced in the rating activity are paper rubric values which will be handed to the chair in the collecting and calculating activity. In a computer-based assessment environment, all values assigned by the rater collected through a question tool will be sent as a response to the collecting and calculating activity automatically. The collecting and calculating activity will use all received rubric values to produce qualification scores, a final score, and the decision according the rules. In the original exam, this task was carried out by the chair manually. Finally, the final score and decision output from the collecting and calculating activity will be used as an input of the informing activity. The artefact will be forwarded to the candidate.

Specifying rules: This sub-section illustrates how to specify the rules for calculating scores and for certificating. We will present the rules in human-readable expressions. First, we have to define properties that represent the rubric values, qualification scores, weights, and decision. Then we specify the rules through defining expressions and actions. In the collecting and calculating activity, the weight of ordinary raters' rubric values and the weight of chair rater's rubric values should be defined. We assume that the weight of an ordinary rater is 1 and the weight of the chair is 2. The weighted average means of the rubric value 'Lead-play and repeat by ear' is calculated using the rubric values of 'Lead-play and repeat by ear (part1)' assigned by three raters:

$$\text{Lead-play and repeat by ear (part1)} = ((\text{rater1.Lead-play and repeat by ear (part1)} + \text{rater2.Lead-play and repeat by ear (part1)}) * \text{rater.weight} + \text{chair.Lead-play and repeat by ear (part1)} * \text{chair.weight})/4$$

For example, if rater1 assigns the 'Lead-play and repeat by ear' for part#1 is 7, rater2 assigns the 'Lead-play and repeat by ear' for part#1 is 8, and the chair assigns the 'Lead-play and repeat by ear' for part#1 is 6, the average means of the rubric value 'Lead-play and repeat by ear' will be: $((7 + 8) * 1 + 6 * 2)/4 = 6.75$. In the same way can calculate the average means of the rubric value 'Improvisation'. The indicator score 'Instrumental-auditive-use' can be calculated as

$$\text{Instrumental-auditive-use (part1)} = (\text{Lead-play and repeat by ear (part1)} + \text{Improvisation (part1)}) / 2$$

A qualification score can be calculated as:

$$\text{Instrumental-auditive-use (total)} = (\text{Instrumental-auditive-use (part1)} + \text{Instrumental-auditive-use (part2)} + \text{Instrumental-auditive-use (part3)}) / 3$$

The final score is the sum of all qualification scores:

Final score = Instrumental-auditive use (total) + Tonal mastery (total) + Technical skills (total) + Rhythmic skills (total) + Ensemble (total) + Expression skills /presentation (total)

The rule for certificating is:

If (final score \geq 36) then (certificate) else (schedule for another exam)

So far we have authored an assessment process model for the use case. We are developing an assessment process authoring tool that will guide the user to develop an assessment process model easily.

10.5 Execution of an Assessment Process Model in an Execution Environment

In order to execute the assessment process model, an execution environment should be developed, which must be able to understand and interpret the assessment process model. As mentioned before, we do not intend to develop a totally new execution environment that is exactly compatible to the assessment process modelling language. Instead, we will design and implement a mapping algorithm to transform an assessment process model to an executable model represented in a combination of LD, QTI, and assessment-specific services. Thus, we can reuse existing LD and QTI compatible run-time environment. The users can benefit from this approach by reusing existing units of learning and QTI resources and seamlessly integrate these with learning and teaching processes.

Taking this approach, the challenge is to transform an assessment process model into a unit of learning with embedded assessment resources. As illustrated in the last section, the assessment process modelling language is designed in the framework of IMS LD. Conceptually it is no problem to translate assessment-specific vocabularies into a generic concept defined in IMS LD. For example, a responding activity can be interpreted as a learning-activity and a rating activity can be interpreted as a supporting-activity. The candidate can be defined as a learner role and the rater is a kind of staff. In a peer assessment, the rater will be mapped to the learner role and the same group of students can be assigned to both candidate role and rater role. Because the transformation algorithm deals with technical knowledge and LD and QTI knowledge in depth, we will not discuss this any detail in this chapter. We will explain how a transformed assessment process model can be executed in an LD+QTI compatible run-time environment.

10.5.1 A LD+QTI Compatible Run-Time Environment

This sub-section introduces a run-time environment for executing unit of learning (UoL) with integrated QTI items. The system architecture is depicted in Fig. 10.7.

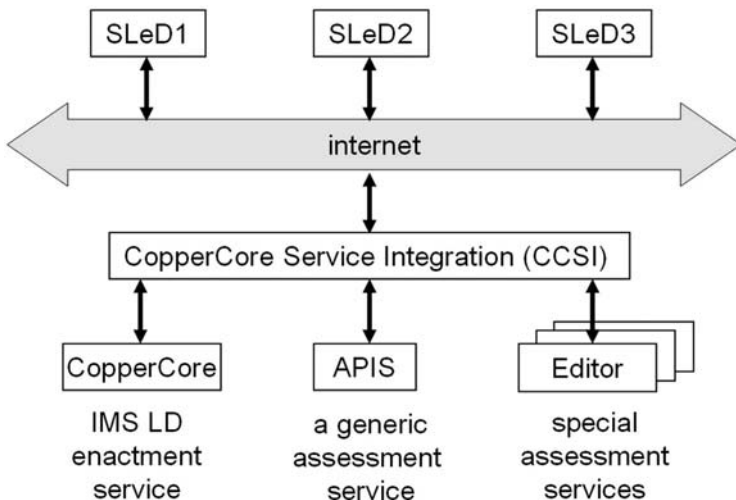


Fig. 10.7 System architecture of an integrated run-time environment

Figure 10.7 depicts an integrated run-time environment, which is implemented by adopting a service-oriented approach. The services include CopperCore (Vogten and Martens 2004) (an IMS LD enactment service), Assessment Provision through Interoperable Segments (APIS) (a QTI run-time service), and some services that can be used to support assessment. Note that the system is open to integrate other purpose-generic services (e.g. whiteboard and text editor) and assessment-specific services (e.g., QTI editor, concept-mapping tool, and LSA tool). Currently, communication to the clients and between services is implemented in a generic integrative service framework, called CopperCore Service Integration (CCSI) (Vogten et al. 2006). The corresponding web based client application is called Service-based Learning Design Player (SLeD).

For users of the system, the client (SLeD, in this example) is the interface to interact with the run-time environment. It is not necessary to know how the services work and cooperate in server side. The next sub-section will illustrate how to interact with the client through using the use case.

10.5.2 Execution of the Use Case in SLeD

The use case presented in Sect. 10.2 can be transformed and wrapped as a UoL. Note that the UoL used here is just a technical term which refers to a digitalized package consisting of a manifest file and necessary resource files. The run-time environment

does not care whether a UoL is a pure learning/teaching model, a pure assessment model, or an integrated teaching, learning and assessment model. Anyway, after publishing the UoL in the run-time environment and creating a run (an execution of the model), the persons involved in the assessment will be assigned to appropriate roles. Many runs of the same UoL can be instantiated, for example, one for each classmate. In the run where Maaïke taking the exam, for instance, the user account for Maaïke will be added in the run. She will be assigned as a candidate. In the same way, three teachers will be included in the run and be assigned as rater1, rater2, and chair, respectively. Note that a teacher can be included in many runs of the same or different UoLs according to the arrangements.

We assume that the exam takes place in a studio equipped with three computers, a large-screen display, and a DVD player. Each rater works with a computer. When Maaïke is ready, the chair starts the run. The DVD player will play the accompanying piano music and Maaïke will play the first part. After Maaïke finishes the first part, each rater will assign rubric values for the first part in the user interface of SLeD. Figure 10.8 shows the user interface of SLeD when the chair was rating the third part.

The left part of the window shows all activities that the chair is involved in. As the icon indicated the chair has finished the first two activities, a currently active activity, and an inactive activity. The chair can shift to other runs by selecting another

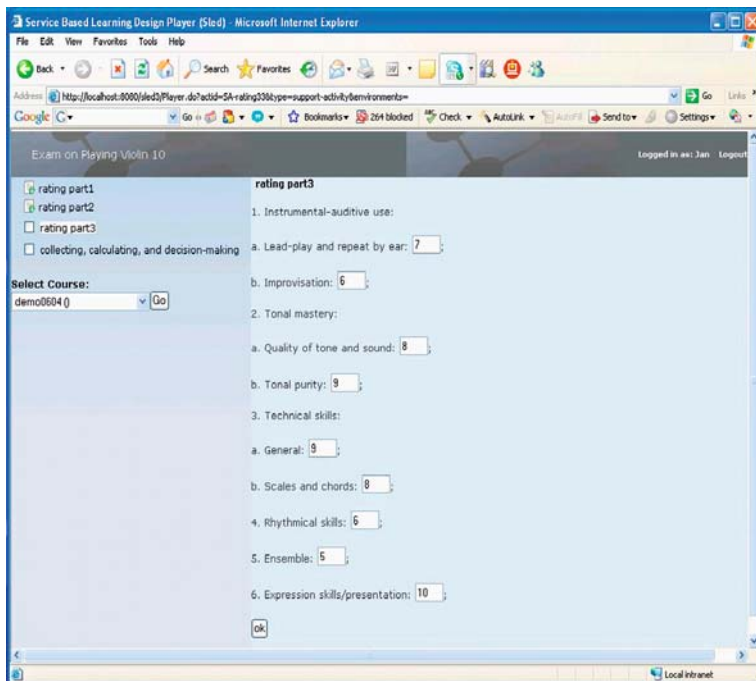


Fig. 10.8 A screenshot of SLeD captured when rater1 was assigning rubric values

course if he/she likes. In the activity ‘rating part3’ The chair is expected to input a number (ranged from 0 to 10) in each blank. After assigning all rubric values, the chair can submit by pressing the ‘OK’ button. After all raters submit their ratings, the exam process proceeds to the next stage, where the activity ‘collecting, calculating, and decision-making’ will be active. The user interface of the chair will change to show all indicator scores and monitor whether all scores are assigned. If everything is in order, the chair can click the small white button in the left of the activity ‘collecting, calculating, and decision-making’ to terminate this activity. As a consequence, the final score will be calculated and the decision will be made according to the definitions of the rules described in Sect. 10.4 Finally, qualification scores, the final score, and the decision will be shown on the large-screen display.

10.6 Conclusion

In this chapter we analysed the limitations of IMS QTI in representing complex assessment, which usually involves multiple roles and participants. The participants perform various kinds of activities in sequence and/or in parallel to produce and consume artefacts in assessment processes. QTI provides insufficient support for modelling control-flow and information-flow in complex assessment processes.

In order to support complex e-assessment in an interoperable and reusable manner, we set out to develop meta-models based on various assessment process models. The first meta-model is the OUNL/CITO model, which provides a conceptual framework for modelling and analysing assessments from a perspective of assessment experts. The second meta-model is the assessment process modelling language, which enables to specify a complex assessment strategy as a formal assessment process model a machine-executable assessment plan. Such a computerised assessment plan can be used as a means to share and exchange assessment experiences and to guide and scaffold learners and other stakeholders in the execution of an assessment within learning networks.

Throughout this chapter we used an authentic exam as an exemplary assessment process to illustrate how to design and describe an assessment process based on the OUNL/CITO meta-model, how to author a formal assessment process model using the assessment process modelling language, and how to execute the formal assessment process model in a run-time environment.

References

- Almond, R. G., Steinberg, L., Mislevy, R. J.: A sample assessment using the four process framework. CSE Report 543 (Center for study of evaluation, Los Angeles, University of California 2001)
- APIS. Assessment Provision through Interoperable Segments: <http://www.elearning.ac.uk/resources/lapisoverview>. Cited May 1 2008
- AQuRate homepage: <http://aquarate.kingston.ac.uk/index.htm>. Cited May 1 2008
- Boud, D., Cohen, R., Sampson, J.: Peer learning and assessment. *Assess. Eval. High. Educ.* **24**, 413–426 (1999)


- Freeman, M., McKenzie, J.: Implementing and evaluating SPARK, a confidential web-based template for self and peer assessment of student teamwork: benefits of evaluating across different subjects. *Br. J. Educ. Technol.* **33**, 553–572 (2002)
- Gehring, E.F.: Electronic peer review and peer grading in computer-science courses. In: *Proceedings of the 32nd SIGCSE Technical Symposium on Computer Science Education*, Charlotte, North Carolina 2001
- Gipps, C.: Socio-cultural perspective on assessment. *Rev. Res. Educ.* **24**, 355–392 (1999)
- Hermans, H. et al.: Educational Model for Assessment, version 1.0 (2006), <http://hdl.handle.net/1820/308>. Cited May 1 2008
- IMS Learning Design specification, version 1.0 (IMS Global Learning Consortium, Inc. 2003). <http://www.imsglobal.org/learningdesign/index.cfm>. Cited May 1 2008
- IMS Question & Test Interoperability, version 2.1 (IMS Global Learning Consortium, Inc. 2006). <http://www.imsglobal.org/question/index.cfm>. Cited May 1 2008
- Joosten-ten Brinke, D. et al.: Modelling assessment for re-use of traditional and new types of assessment. *Comp. Hum. Behav.* **23**, 2721–2741 (2007)
- Koper, E. J. R., Olivier, B.: Representing the learning design of units of learning. *Educ. Technol. Soc.* **7**, 97–111 (2004)
- Lockyer, J.: Multisource feedback in the assessment of physician competencies. *J. Contin. Educ. Health Prof.* **23**, 4–12 (2003)
- Miao, Y., Sloep, P., Koper, R.: Modelling units of assessment for sharing assessment process information: towards an assessment process specification, In: *Advances in Web Based Learning – Proceedings of the 7th International Conference on Web-based Learning*, ed. by Li, F. B. J et al., ICWL 2008, Jinhua, China, August 2008. *Lecture Notes in Computer Science*, vol. 5145 (Springer, Berlin Heidelberg New York 2008a) pp. 132–144
- Miao, Y. et al.: Bridging the gap between practitioners and e-learning standards: a domain-specific modelling approach. In: *Times of Convergence. Technologies Across Learning Contexts – Proceedings of the Third European Conference on Technology Enhanced Learning*, ed. by Dillenbourg, P., Specht, M., EC-TEL 2008, Maastricht, The Netherlands, September 2008. *Lecture Notes in Computer Science*, vol. 5192 (Springer, Berlin Heidelberg New York 2008b) pp. 284–289
- SLeD: Service Based Learning Design System (2005) <http://sled.open.ac.uk/>. Cited May 1 2008
- Topping, K. J.: Peer assessment between students in colleges and universities *Rev. Educat. Res.* **68**, 249–276 (1998)
- Vogten, H., Martens, H.: ‘CopperCore’ <http://www.coppercore.org>. (2004)
- Vogten, H. et al.: CopperCore Service Integration – integrating IMS Learning Design and IMS Question and Test Interoperability. In: *Proceedings of the 6th IEEE International Conference on Advanced Learning Technologies*, Kerkrade, The Netherlands (IEEE Computer Society, 2006) pp. 378–382
- Volder, M. D. et al.: Espace: a new Web-tool for peer assessment with in-built feedback quality system. In: *Proceedings of ABR & TLC Conference, 2007*, Hawaii, USA, 2007
- Wills, G. et al.: R2Q2: rendering and responses processing for QTIv2 question types. In: *Proceedings of the 10th International Computer Assisted Assessment Conference*, ed. by Danson, M. (Loughborough University, UK 2006) pp. 515–522
- Zoho Challenge: <http://zohochallenge.com/OnlineTest/Candidate.do>. Cited May 1 2008

Chapter 11

Placement Services for Learning Networks

Jan van Bruggen, Marco Kalz, and Desirée Joosten-Ten Brinke

Learning



11.1 Introduction

Placement in Learning Networks, or positioning as we have called it elsewhere (Van Bruggen et al. 2004), is a process by which we try to put a student on the most appropriate place in one of the alternative paths leading to a specific competence, taking into account the goals and preferences, as well as the history of the professional.

Placement, although related is different from navigation which is described elsewhere in this volume. In terms of the terminology presented above, navigation assists the professional in *making decisions within* a particular path. Navigation assists the student in making a decision on which learning activity to engage in next (when this is an option at all). Placement, in contrast, seeks to identify an effi-

J. van Bruggen (✉)
Centre for Learning Sciences and Technologies, Open University of the Netherlands,
Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands
e-mail: jan.vanbruggen@ou.nl

cient path towards a goal by identifying those learning activities in one or more paths that the student might skip, taking into account the professional history. Navigation is dealt with in more detail elsewhere in this volume and we refer the reader to the chapters in that section for more elaborate descriptions of navigation services offered in Learning Networks.

Placement in Learning Networks is complicated by two factors:

1. The potentially very large number of contributions to the learning network combined with the absence of a common or controlled vocabulary in the description of learning activities and their outcomes, and
2. The guideline for professional development to not only consider the professional's history in formal learning, but to consider *all* prior learning for credit.

We discuss both complicating factors in the next sections.

11.2 Contributions to the Learning Network

The open nature of the environment offered by Learning Networks complicates placement, in particular its implication that all members of a learning network may freely contribute learning activities or even complete curricula to the network. A learning network may grow to contain a wide range of opportunities to build or strengthen professional competencies in the domain targeted by the learning network. Thus, a new member is confronted with a potentially large set of educational offerings that only share a common denominator in the domain that is being addressed. Apart from that common core, there may be offerings ranging from accredited curricula provided by educational institutions that lead to a formal, full degree, to collections of materials of non-formal nature.

A road towards a specified competence goal may be predefined, as is the case in a curriculum, whether it is offered by a formal institution (a major in Music Performance offered by a college of music), or in a non-formal way by a single actor ("Doctor Guitar's road to the stage"). It is likely that such a learning network will expand to offer various roads leading to a competence outcome their contributors claim to be the same or much alike. Several predefined roads towards guitar performance, some of them leading to an accredited degree, will become available. We refer to these defined roads as "learning paths".

Next to complete learning paths towards a predefined competence goal, what we called a curriculum, there may be a vast amounts of resources, not bound together in a curriculum, that are contributed by individual actors. Today's Internet already offers a plethora of lessons, examples, transcriptions et cetera that are made available by large numbers of amateur and professional guitar players. In a learning network there will be overlaps in the learning activities contained in these predefined roads: several universities will offer master programmes in Psychology; different actors, including institutes for formal music education, will offer roads leading to

professional guitar playing. And, as already witnessed in the World Wide Web today, there will be vast amounts of non-formal learning resources that are offered, on an as-is basis by various individual or collective actors.

All this makes placement a process in which extended, but not necessarily complicated, search through alternatives is necessary. The real complicating factor is the absence of a common vocabulary. A learning network does not enforce its users to adopt a common controlled vocabulary to describe the learning activities and its learning outcomes. Note that the latter also has implications for the portfolio of members of a learning network.

Although, as emphasized above, no controlled vocabulary to describe competencies and learning activities is enforced in Learning Networks, this does not preclude that areas with a common vocabulary exist. Metadata on curricula offered by European institutes may conform to the Dublin descriptors or to the European Qualification Framework. Even in these cases it can be challenging to recognize that metadata are compliant with a controlled vocabulary and to identify that vocabulary. It seems likely that compliant parts of the learning network are associated to providers offering formal education. The real challenge for placement then is to use *whatever data is available* to locate material and to map that onto the learner data available.

11.3 Placement and Accreditation of Prior Learning

Placement is one of the traditional purposes of assessment, and it has become associated with decisions on (conditional) admission to a learning program, or on allocations to jobs and functions. In Learning Networks professionals select curricula or separate learning activities that matches their development needs, goals and preferences. A placement service should assist the professional in selecting appropriate material and, in doing so, the service should also take into account the professional history. In the paradigm of professional development that is underlying Learning Networks that means that *all* prior learning, irrespective of its nature or source, should be considered for credit. Only in the rare case that a professional's history and current activities share the descriptors for competencies and content, can placement decisions be based on evidence gathered through assessment procedures that are fit to measure competence development. Whenever a professional crosses boundaries between regions that do not share a common description, placement decisions or advice, are based on what is called Accreditation or Assessment of Prior Learning (APL). Now the professional is seeking exemptions for obligations on the basis of prior learning, regardless the context in which this learning occurred. Placement in Learning Networks can be understood as a process of Accreditation of Prior Learning (APL).

As we emphasized, all forms of prior learning, or rather their outcomes are considered in APL procedures, ranging from formal education where achievements are related to recognized certification to professional experience. All types of learning

can be taken into account (as a source) in an APL process. Placement in Learning Networks works in the same way.

Joosten-Ten Brinke et al. (2008) provide an overview of a number of characteristics of APL procedures and discuss quality aspects of this type of assessment. Of particular interest to the topic of placement in Learning Networks is Joosten-Ten Brinke et al. (submitted). They used groups of experts and students to identify services and the estimated added value and efficiency of each of them. The services were divided into embedded support, that is support contained in printed or electronic materials and personal support that is offered face-to-face or (a)synchronously. From the results the authors identified a framework that links the phases of APL to the services that offer added value and efficiency. Obviously for professional development in Learning Networks any heavy reliance on human support in the APL process will be difficult to maintain and we will therefore concentrate on the embedded support identified.

The following phases in the APL process are encountered commonly in the literature:

- 1 In the *profiling* phase, the institution gathers information about the student's personal characteristics and needs. In this phase, the institution can inform the student about the steps and expectations of the procedure. The suggested embedded support in this phase consists of at least an APL manual, a self-assessment instrument to test whether the procedure is likely to be meaningful, a website with FAQs and information about APL, and finally good and bad examples of portfolios with clarification. Of particular interest here is the self-assessment instrument that could be offered as one of the services for placement in Learning Networks.
- 2 In the phase of *gathering and presenting* the evidence students collect and present evidence about their qualifications and experience to support a claim for credit with respect to the new qualification they are seeking. The burden of proof is on the students: they have to provide evidence for their claim that they have acquired knowledge, skills or competences that meet the requirements of the course or learning program that they wish to follow. The evidence is typically collected in a portfolio. This is the phase that is the most troublesome in the whole process where both students and institutions require assistance. The critical phase is the second one, where students have to gather and present evidence. Students report serious problems with the preparation of a portfolio that demonstrates the equivalence of prior learning to the outcomes of the target courses. This student problem is, interesting enough, mirrored by the institutions; they seek assistance in mapping the data that the students provide in their portfolios onto their courses and the expected learning outcomes. The embedded support suggested here includes a "mind-manager system with a portfolio format" and strong versus poor examples; a "how to compose a portfolio" manual; the opportunity to electronically seek and present analogous cases, and a FAQs list. The system alluded to here would draw the attention of the learners to slots to be filled in the portfolio structure. A more daring interpretation is that it would allow an external representation

where content in the portfolio – which could contain a diversity of data, such as certificates, course descriptions, products produced by the student – is mapped on the outcomes of the target learning activities.

- 3 In the phase of *assessing the evidence*, the quality of the student’s evidence is reviewed by assessors confirming to assessment standards. The result of the assessment is an answer to the question whether the student should gain recognition. Embedded support can consist of a list of assessment criteria, an elaboration of the deployed assessment protocol, examples of good and bad portfolios for competence assessment, and an overview of assessment results jurisprudence. We will not consider most of these embedded support options any further since they are only providing information about the APL process to the learners. Note that in the placement service we envision, the jurisprudence as well as the examples are part of the case-based reasoning part of the service.
- 4 In the final phase of *accreditation* (or “recognition”) these results are verified or endorsed by the department responsible for awarding the credit or recognizing the outcome of the assessment. Embedded support suggested by the authors should include examples in which recognition was and was not given, descriptions of standard recognitions and the phase itself, and graphic overviews of the educational programme.

Some of the steps in the APL process can be framed in terms of the model of assessment proposed by Pellegrino et al. (2001). Their model of assessment has “Cognition”, “Observation” and “Interpretation” as main elements. Cognition is a model of how a learner represents knowledge and develops competencies; Observations are tasks or situations in which (complex) behaviour can be observed, and interpretation is a means by which one can make sense of the observations. Several examples of new “linkages” between the main elements are given by the authors, such as the use of concept mapping to assess knowledge structures (linking Cognition to Observation), or the use of Latent Semantic Analysis to interpret essays (linking Observation to Interpretation). In this model items (tasks) provide part of the evidence that is linked to the learning objective and must support decisions that are based on the assessment results. These items selected for observation should be developed with the purpose of the assessment in mind (i.e. going from cognition to observation). The evidence gathered still needs to be interpreted and this interpretation expresses how the observations constitute evidence about the learner’s knowledge and skills.

In contrast to this assessment model, APL procedures in general cannot use such constructed linkages that allow making inferences about the learner’s competence in a straightforward way. Rather, the analysis of the portfolio is an exercise in interpreting its content as proxies to elements of the assessment model for the target. Consider the following cases:

- (α) The portfolio contains certificates of courses that entitle students to exemptions of target courses. In this case the certificates are directly *taken* as evidence of the target outcomes and no further interpretation is required.

- (β) The portfolio contains competence maps of prior learning by the student that one may attempt to match onto similar descriptions of the target learning activities or curriculum, basically trying to interpret the student's data as matching the Cognition aspects. Certificates of non-recognized courses will be processed in a similar way, which is through descriptions of the course outcomes.
- (χ) The portfolio contains products produced by the student. Now the matching takes two steps. First, the linkage between (target) Cognition and Observation is handled, that is one assesses whether the type of product (an observation) is a valid demonstration of the target Cognition. Second, the content of the product itself is inspected, that means it is interpreted to determine whether it contributes to the evidence for the target Cognition. Interpretation of observations can be complex, for example when one has to match material that is phrased in different terms than the target domain (Starr-Glass 2002; Starr-Glass and Schwartzbaum 2003).
- (δ) The portfolio contains (links to inspectable) contents studied by the student and possibly certified in a non-formal manner. Here one can compare student content to that of the target domain as a (rather remote) proxy of the outcomes. The core assumption here is that processing similar content will lead to similar outcomes.

Figure 11.1 relates the APL procedures of portfolio inspection to the assessment triangle of Pellegrino et al. (2001). In the figure we have only included the comparison between student competence maps to target competence maps, but note that in

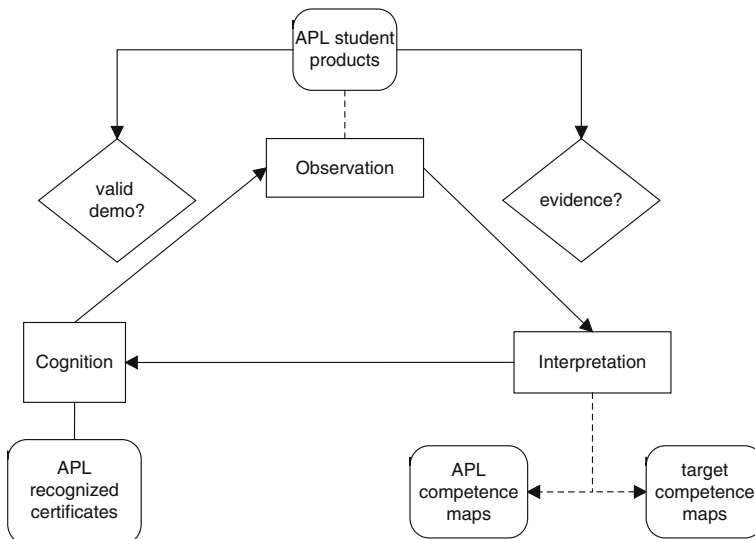


Fig. 11.1 APL procedures linked to the assessment triangle

order to interpret student products, content or certificates of non-recognized courses we need similar descriptions of the target courses or curricula. This observation has led to what we call “symmetrical positioning”. Our approach to placement in Learning Networks has been influenced by the absence of controlled vocabulary in Learning Networks, which limits the data available for placement decisions to cases (c) and (d). This work is described in somewhat more detail in the next section.

11.4 Design of Placement Services in Learning Networks

The design of services for placement are based on a number of design principles.

First, one of the leading design principles in the development of services for assessment and placement has been that the involvement of human experts in preparing placement decisions should be kept to a minimum. We therefore concentrated on computer-based placement services

Second, the placement services are designed to inform and advise professionals on their placement in the learning paths considered. That is they can be used to inspect the portfolio and identify strong and weak contents for the professional. Obviously, the same functionality can be used for the provider.

Third, placement services respect the distinction between the assessment of the portfolio and the endorsement. The placement services present the results of the analysis along with the consequences – credits and, exemptions for particular learning activities. How the analysis and its consequences are used, depends on the policies of the provider. In the most restricted of policies the use of the placement service could be limited to providing input to human assessors. In a far-reaching scenario the policies of the provider would consider the analysis and consequences sufficient for automatic endorsement. However, the more likely case in Learning Networks is that the services are used to assist assessors in an APL procedure in the evaluation of the professional’s portfolio, and to give some idea about potential credits. Within the framework of formal education, policies with respect to placement would most likely preclude any placement decisions by external services.

Finally, services must operate under a minimum set of assumptions. As our remarks above indicated, within Learning Networks one cannot assume a common vocabulary. In APL procedures the description of the learning outcomes of the target is normally known and remains fixed (although several different targets can be considered). One then seeks to relate the contents of the portfolio to one or more targets as we described above using the assessment triangle of Pellegrino et al. (2001). In Learning Networks, however, neither source nor target can be assumed to be fixed. Under these data constraints we have to consider what type of data could be available to describe the target learning activities on the hand and the contents of the portfolio on the other. We consider three cases, corresponding to different types of data: informal descriptions, metadata descriptions, and ontological descriptions.

We use the term “informal descriptions” to denote the case where the professional enters without any explicit competence descriptions in the portfolio. Competence development programmes with informal descriptions offer no information in a structured manner about the competence resulting from the learning activities. As an example, consider a professional who wants to update his competence in accounting. He enters a learning network that deals with the domain of finance and accounting by offering a number of learning activities. The learner feels his competence development goal will be reached by working through a collection of the learning activities offered. The electronic portfolio of the professional contains only some of the documents he produced in his former education. After completing the learning activities the portfolio will now include products produced in the finance and accounting learning activities as well as pointers to these learning activities and their content.

The second case is one in which metadata are used to describe the professional competencies as well as the contents and outcomes of the learning activities. For example, if a professional enters with a standards-compliant ePortfolio the situation for the placement service would be different. Consider the example of case one again. This time there are metadata; so now the professional enters with a standards-based description of his (entry) competencies and the learning activities in the learning network contain detailed and standardized information about their entry requirements and their competence outcomes. Elsewhere we have described several existing schemes that offer interoperable metadata for competences and portfolio, including the IMS-based reusable definition of competency or educational objective (RDCEO) specification and the HR XML Consortium competency schema for a variety of business contexts that is applicable in recruitment processes (HR XML Consortium 2004). Although these metadata specifications may provide reusable machine-readable metadata, their support to placement decisions is limited to those cases in which matching characteristics of the contents of the portfolio to target learning activities and outcomes is rather straightforward.

The third case is one in which learner profiles and learning activities are described in an ontology. Ontologies are metadata schemas providing a controlled vocabulary of concepts; They can be useful in sharing a common understanding in a domain in a machine-readable way. For competence development ontologies or taxonomies can be used to define competencies related to learning activities. Competence ontologies could be added to either the student profiles (Dolog and Schaefer 2005), learning objects (Ng et al. 2006) or the competence development programmes (Woelk 2002). An ideal ontology encompasses the contents and content structure of a domain, as well as the learning outcomes and how they can be achieved. Thus, the ontology combines a specification of the domain with a specification of the development of professional knowledge, skills and competence in the domain. Placement inside a learning network can be based on the relations between domain ontology and the competence ontology (Posea and Harzallah 2004). The process of adding competencies to the student profile could proceed from successfully finished learning activities in the learning network. Parts of the competence

ontology in the learning network could be added to the student profile step-by-step after they have successfully completed the related assignments.

Whereas metadata allow comparison and matches on descriptors, ontologies allow advanced reasoning as well. For example, if we observe a student who demonstrates mastery of advanced data analysis techniques, we may infer that this person will master simple techniques as well. If someone produces a text on a specialized topic in psychology we infer (as did the experts in one of our experiments in placement) that this person possesses knowledge of a general introduction to psychology. In this way ontology allows to infer beliefs about the knowledge, skills and competences of the student.

This is the type of reasoning that human experts use when they interpret the contents of a portfolio in terms of a target curriculum. Consider the reasoning about a product as depicted in Fig. 11.1. Assume that the product is part of a portfolio and that the target – for which the student is demanding an exemption – is an introductory course in psychology. Let us assume further that the portfolio contains several learner written papers dealing with topics in psychology. Experts readily recognize whether the topics are relevant and they know what the intricacies of these topics are. They also recognize whether the papers treat the topics in a basic or in a more advanced manner and taking all these together they might conclude that the student is already performing beyond the level of an introductory course and so decide to exempt this course.

The symmetrical positioning approach referred to above, results from considering the types of data that can be available on the student and on the learning network and applying the simplifying assumption that we concentrate on the symmetrical case were these descriptions are cast in the same type of data. This is depicted in Table 11.1.

We also decided to first address the situation where only informal data are available. These are the cases (c) and (d) in the description of APL procedures in placement, that is a situation in which we have access to contents studied and products created by the professional and perhaps non-structured associated descriptions. We have referred to placement under these conditions as “content-based positioning”. The rationale and the research agenda for this approach to placement has been described in Van Bruggen et al. (2004). The approach rests on the following assumptions. In general there will be an absence of the extensive assessments

Table 11.1 data for symmetrical positioning

–	Structured data	–	–	–	–
Learner information	Ontology	–	–	X	–
	Metadata	–	X	–	–
	Informal	X	–	–	–
	Unstructured data	Informal	Metadata	Ontology	Structured data
	Learning network				

that can directly demonstrate that the professional has already acquired knowledge, skills and competences that are equivalent to the outcomes of learning activities that are being considered. Note that the Miao and Hermans (Chap. 10) in this volume details the design of a model that supports assessment types needed for competence assessment, but does not address the issue of unification of assessment descriptions in the learning network. Thus, the main data source on which placement decisions can be based are informal descriptions of competences and material as well as the learning material itself. The core assumption is that equivalence of outcomes can be approximated by the similarity of the contents of (learning) materials studied or produced by the student (source material) and the material contained in the learning activities in the learning network (target). If a placement service determines that the content of source and target materials overlap substantially, we infer that the professional has already acquired the outcomes of that activity.

In our work on content-based placement document similarity is computed using techniques inspired by latent semantic analysis (LSA) (Deerwester et al. 1990). In LSA documents are represented as a vector with term frequencies and brought together with other document vectors in a so called term-document matrix. Using a technique comparable to principle components analysis the dimensions of the matrix are determined and a dimension reduction of the matrix is performed by dropping the smaller dimensions.

Professionals are represented by one or more documents that they have produced or studied. If one or more of these document vectors show a high correlation with learning material vectors, then the learning material may be considered redundant. Although the content-based approach has modest requirements on the way data are expressed, there are several limitations and assumptions that we need to consider, such as the amount and quality of available material in the profiles for content-based placement.

Although placement in Learning Networks is not a topic addressed in LSA research and development, there are several examples of related research that show promising results. Zampa and Lemaire (2002) modelled a domain as a set of topical contents (lexemes) and used LSA to compute a “Zone of Proximal Development” from which these topics were selected and sequenced. This is content-based placement on the scale of a learning activity, but it needs hand-crafted “lexemes”. On a different scale LSA was used to identify and relate people, jobs and training. On the basis of job descriptions for three occupations as well as “duty lists” and actual tasks performed Laham et al. (2000) created a semantic space for jobs and people. The semantic similarities between jobs and people could be used to decide between candidates for the job or to select a replacement.

11.5 Methodological Considerations

All placement methods, whether they are human-based, computer-based or mixed, have to adhere to the qualities of reliability and validity. Any placement service should work in a reliable manner, that is similar cases should be treated equally

within an acceptable error margin. The maximum reliability for machine-based methods is dictated by those of human experts – we will not expect computer based placement to outperform human experts who will not always agree. A reliable operation of a placement service is a precondition for its validity. The validity of a placement service refers to the objective correctness of the placement decision. In our approach this has always been interpreted as a performance comparable to that of a human expert.

The core assumption of the placement work reported is that if material studied or produced by a new member of a learning network correlates significantly with material contained in the target learning activities then there is evidence that this student has already acquired the learning outcomes associated with the learning activities. Thus, we consider material, in particular textual material, as a proxy to competence development. Whenever material correlates significantly this is considered as an indicator that they “tap” the same competence.

The technique used to compute these correlations is related to latent semantic analysis. More properly, it should be termed “reduced vector space model”, because its modelling is based on an initial vector model of documents whose dimensions are reduced, and does not use the large corpora used in LSA to train the system on the general language. Corpora of written Dutch of the size needed to train LSA are not available. Kalz reports results that used one of the small sized corpora (Parole distributable) that are available, but found no difference after training with this corpus. In these models documents are represented as a single vector of (transformed) term frequencies. These vectors can be correlated as they are (raw) or their projections in a reduced space environment can be correlated, in a procedure that is closely related to principal components analysis. The calculation of reduced vector space models is not a straightforward process, but requires selection and often optimization of a broad set of parameters ranging from decisions on the smallest size of terms to include, on the type of frequency measure to use to the number of dimensions to retain in the final model. Although a number of guidelines have more or less emerged from research of laboratory experience there is no recipe available. We will not go into any detail about the details of these calculations since these are discussed in Chaps. 12 and 13 of this volume.

There are several weaknesses and caveats connected to the type of content-based placement discussed here.

First, content-based placement makes placement decisions based on the correlation between portfolio documents and a single target document, that is, it tests a one-to-one correspondence. This may not be too unrealistic in comparison to APL procedures but it will return weaker results if two or three learner documents combined would match a target document. In theory it is quite doable to test different combinations of documents since the text vector model is additive, but the rules for this type of engineering of “re-combinations” of documents are not obvious.

Second, the model cannot use a hierarchy of competences to professional competencies. Take the example of a text that describes an advanced experimental setup to test a hypothesis by describing tasks, measurements, and data analysis. From this description any expert in the domain would infer that the writer masters content

and skills at a proficient level. Nevertheless, when we correlate such a text with the standard introductory texts we will find only modest correlations. In all likelihood such a text would produce a higher correlation when compared to an advanced text in experimental research. Unfortunately, as long as learning activities can only be interpreted as a string of activities without any clue as to the corresponding level of competence, there is no way for a machine-based service to correctly interpret the input, that is to make the appropriate inference regarding the actual level of professional competence in the domain.

Third, correlations between documents are normally based on the cosine between vectors, which is a measure in which the length of the texts is not considered. Thus, at least in theory, a small source text could be successfully matched against a much longer target text like “Introductory Psychology”. A comparison of the lengths of the vectors involved could be added as a second factor in placement services.

A further limitation is closely related to the accuracy and flexibility with which experts use the concepts in the domain. Although experts and novices tend to use the same terms to describe concepts, novices – when asked to compare and contrast these terms and concepts – typically come up with scattered conceptual models whereas experts produce alternative conceptual structures that correspond to different views on the domain or a problem. In their description of individual terms experts are more precise and articulate. This is not evidenced by the use of technical or domain-specific terms but rather by the use of vague terms by novices – “this relates to...”, “this leads to” – in contrast to experts who provide more specific descriptions of the relations. Whether the technique can discriminate this usage remains to be seen.

Finally, we may encounter a situation in which the target material itself does not allow sufficient discrimination. This seems to have been the case in the example of “Introductory Psychology” that will be presented in Chap. 13. The large communality in the chapters of this introductory course are reflected in correlations that are so high that they prohibit discrimination between the chapters, at least they did with the traditional vector correlation techniques that we used. One of the lessons learned there is that we have to ensure that the service has sufficient sensitivity. This issue will be dealt with in Chap. 13.

The limitations also point to directions in which to improve the placement services. Here we concentrate on suggestions to focus on particular content or to extract other types of content that is in the material. A first approach is to extract information that is in particular salient to discriminate between learners at different competence levels. As the data presented above indicate, making a successful discrimination between experts and novices – with text vector models – may need a filtering on *modal terms* rather than on concepts. This, however, will not help in a case where the material contained in different learning activities cannot be discriminated. One of the suggestions done in the literature is to use summaries or lists of important concepts here.

11.6 Conclusion

Placement of a professional in a learning network for professional development is a complex task in itself and this is exacerbated by conditions that prevent any simple mapping of profiles and competency descriptions onto the educational resources. The two most extreme situations that we considered are the clearest: (1) no competency descriptions inside the profiles and the learning network; (2) competence ontologies in the profile and the learning network. In the first case a content-based approach is the one to take. The content-based approach to the placement problem has the advantage that it can be used for placement here and now, where most professionals will not have explicit competence descriptions in their profiles. The drawback of the approach is that it is only related to the documentation produced by the professional learner and not to his acquired competencies. So success is dependent on the documentation that is provided in relation to the educational and professional history. If, for example there is only content for parts of the educational and professional background, the placement recommendation will be biased. In addition, the concentration on content may effectively limit the approach to domains with a strong verbal character. For the same reason, domains with psycho-motor content, for example practical skills, may not be adequately represented.

References

- Deerwester, S. et al.: Indexing by latent semantic analysis. *J. Am. Soc. Inform. Sci.* **41**, 391–407 (1990)
- Dolog, P., Schaefer, M.: A framework for browsing and manipulating and maintaining interoperable learner profiles. In *Proceedings of UM2005, 10th International Conference on User Modeling*, Edinburgh, UK, July 2005 (Springer, Berlin, Heidelberg, New York 2005)
- HR XML Consortium: *Competencies (Measurable Characteristics)* 2004. http://ns.hr-xml.org/2_3/HR-XML-2_3/CPO/Competencies.html. Cited July 1 2006
- Joosten-Ten Brinke, D. et al.: The quality of procedures to assess and credit prior learning: Implications for design. *Educ. Res. Rev.* **1**, 51–65 (2008)
- Laham, D., Bennett, W., Landauer, T. K.: An LSA-based software tool for matching jobs, people, and instruction. *Interact. Learn. Environ.* **8**, 171–185 (2000)
- Ng, A., Hatala, M., Gasevic, D.: Ontology-based approach to learning object formalization. In *Competencies in Organizational E-Learning: Concepts and Tools*, ed. by Sicilia, M. -A. (Idea Publishing, Hershey 2006) pp. 185–206
- Pellegrino, J. W., Chudowsky, N., Glaser, R. (eds.): *Knowing What Students Know: The Science and Design of Educational Assessment* (National Academy Press, Washington, DC 2001)
- Posea, V., Harzallah, M.: Building an ontology of competencies. *Proceeding of Workshop on Ontology and Enterprise Modelling: Ingredients for Interoperability*. Fifth International Conference on Practical Aspects of Knowledge Management. Vienna, December 2004.
- Starr-Glass, D.: Metaphor and totem: exploring and evaluating prior experiential learning. *Assess. Eval. High. Educ.* **27**, 221–231 (2002)

- Starr-Glass, D., Schwartzbaum, A.: Liminal space: challenges and opportunities in accreditation of prior learning in Judaic Studies. *Assess. Eval. High. Educ.* **28**, 179–192 (2003)
- Van Bruggen, J. et al.: Latent semantic analysis as a tool for learner positioning in learning networks for lifelong learning. *Br. J. Educ. Technol.* **35**, 729–738 (2004)
- Woelk, D.: E-learning, semantic web services and competency ontologies, *Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications*, 2002.
- Zampa, V., Lemaire, B.: Latent semantic analysis for user modelling. *J. Intell. Inf. Syst.* **18**, 15–30 (2002)

Chapter 12

Tools and Techniques for Placement Experiments

Wim van der Vegt, Marco Kalz, Bas Giesbers, Fridolin Wild, and Jan van Bruggen



12.1 Placement Support with Language Technology: Probabilistic Reasoning About Prior Knowledge Based on Semantic Analysis

In Chap. 11 we presented placement in the context of Accreditation of Prior Learning and showed that in the scenario we address we do not assume the availability of controlled vocabulary with which the contents of the learner portfolio or the learning material in the learning network is described. Our placement service is based on the assumption that similarity between material produced or studied by the learner on the one hand and the learning material of the learning network on the other, can be used as a proxy to similarity in learning outcomes. The first task of any such

W. van der Vegt (✉)
Centre for Learning Sciences and Technologies, Open University of the Netherlands,
Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands
e-mail: wim.vandervegt@ou.nl

placement service is therefore to establish whether these similarities are present for a given learner. The technology with which this is done, latent semantic analysis, is presented here. The emphasis here is on the technical and computational aspects of data preparation and analysis. In the next chapter two cases are presented in a framework for validation of a placement service where the emphasis is on analyzing and interpreting the data in order to verify that a placement service is reliable and valid.

In general, the use of a method like LSA to support the placement process is always a probabilistic method. There are two main reasons, why this is the case. First of all, the similarity calculation between concepts and documents is based on heuristics expressed in the deployed document collection. With a different corpus, varying similarities may be achieved. To a certain degree, these differences can be balanced out through carefully selecting the training corpus used by the algorithm. Second and not less important, the calculated similarity of concepts in a domain expressed in the portfolio documents and in the targeted course content is itself only a probability as it assumes that content similarity can be used as an abduction rule for assessing prior knowledge. This assumption, however, is not far fetched as it is very likely that domain experts use the same mechanism to decide on prior knowledge and study exemptions. In our current work we are focusing on semantic similarity while there are also several other possible methods to capture prior knowledge from more structured data (Kalz et al. 2007).

Latent Semantic Analysis (LSA) is a theory and method for extracting and representing the meaning of words by statistical computations (Landauer et al. 1998). It provides a method to calculate the similarity of texts and parts of texts. The process of analysis consists of several steps, most prominently pre-processing, weighting and normalizing mechanisms, the construction of a Term-Document matrix, and, – at the core – the singular-value decomposition (SVD), a mathematical factorization method similar to principal component analysis (PCA). The end result of this process is a rank-reduced latent-semantic space, in which the main concepts (aka “types”) and the documents of the input are represented as *vectors*. Concepts in this space are similar if they appeared in the same contexts: then their vectors are close, as the constituting values along each dimension (i.e., singular value) are similar. This way, the distances between vectors provide a measure for the similarity of the concepts they represent. The angle, for example, between the vectors representing two similar concepts will be small, whereas the angle between two unrelated term vectors will be big. LSA is applied in several research fields like informatics, psychology, or medicine.

For technology-enhanced learning, the application of Latent Semantic Analysis may alleviate problems such as increased tutor load, or the provision of formative feedback during learning (Van Rosmalen et al. 2006; Graesser et al. 2000). Some of the most prominent examples for the use of LSA in an educational environment are the assessment of and feedback on free-text answers in intelligent tutoring systems. Examples of these applications are the Intelligent Essay Assessor (Foltz et al. 1999), Summary Street (Steinhart 2001), and Select-a-Kibitzer (Wiemer-Hastings

and Graesser 2000). Others have used LSA to provide students with text that is appropriate to their current knowledge (Wolfe et al. 1998; Dessus 2004).

Using LSA for placement in Learning Networks is similar, but has a different motivation and context. LSA is being used to assess prior knowledge of learners for placement, to construct personalized learning paths, and to create individualized curricula through a learning network. In this chapter we concentrate on the data and the tools for LSA. In the next chapter the validation of a placement service is discussed.

12.2 Corpora and Datasets for Placement Experiments

For every placement experiment based on LSA a set of text collections is needed. Although there are many text corpora available from the wider field of Natural Language Processing (NLP), most of them cannot be used as-is for experimental research with LSA: many corpora are not freely accessible; the available corpora and datasets are heavily annotated with e.g. Part-of-Speech-Tags (POS-Tags); or a corpus matching the targeted domain may turn out to be too small. Corpora needed for Latent Semantic Analysis contain a considerable amount of raw texts without any mark-up. Documents in such a corpus are treated as so-called “bag-of-words”, i.e. regardless of their word order while respecting their context. When we talk about corpora and datasets for LSA we can distinguish between three different types:

- A large set of documents that define the underlying “language” (the “background” corpus).
- A smaller set of documents that define the domain of interest in which the analysis takes place, including, for example, the content of the study program in which the experiment takes place (the “domain-specific” corpus).
- One or more query documents which are to be analyzed by investigating their similarity relations (the “test” corpus).

The language corpus consists of a large set of documents that can represent the use of semantics in the chosen language. Quesada (2007) formulates an ideal corpus for LSA as follows: “One has a corpus big enough when, for any new learning experience added, the probability of adding a new type is so low that it is negligible”. Landauer (2007) reports good experiences with general language corpora containing 10^5 to 10^{10} words of text divided into 10^5 to 10^9 paragraphs with 10^5 to 10^6 different word types. However, it is not always clear what is meant by “large” and “small” corpora, or what size is perceived as a minimum and/or maximum requirement to successfully apply LSA. LSA is often used for document retrieval from very large document databases, containing ten thousands of documents and an input of 5000 documents to train LSA on the domain is quite common in these applications. Deerwester et al. (1990) state that a “reasonable size” is 1000–2000 abstracts which means about 5000–7000 index terms. In contrast, in Laham et al. (2000) a total

number of 20,000 objects is mentioned as a very small dataset compared to LSA's capabilities, but enough to do a fair job in estimating statistical regularities. Many authors seem to take this as a minimum requirement.

Educational uses of LSA, in contrast, are often confined to smaller corpora, that are more specific to (sub)domains. Within these corpora LSA has been shown to be robust against decreasing the size of the corpus: According to Wiemer-Hastings and Graesser (2000), the best corpus is specific enough to allow for subtle semantic distinctions within a domain, but is general enough to ensure moderate variations in terminology won't be lost. When they reduced the size of their text corpus from 2.3 MB to a minimum of 15%, the effectiveness of LSA in terms of correspondence with human raters decreased 12%. Obtaining meaningful LSA results from smaller corpora is possible as shown in Wild et al. (2005) where 43 files each consisting of a student's answer on a marketing question from a real world exam were analyzed and showed significant correlation between machine scores and human scores. Also a corpus used for sensitivity testing (see Chap. 13) produced meaningful results, small as it was, containing less than 300 documents. When working with small corpora, the generic "language" background corpus is omitted and only a domain-specific corpus is used. In this case, the domain-specific corpus has to serve both purposes: to train the semantic structure of the target language used and to teach the particularities of the domain of interest. This approach only seems to work under certain restrictions and easily degrades effectively below the performance of the pure vector space model (not only with unfavorable parameter settings). Therefore, one would normally prefer to use corpora that are considerably larger.

If researchers want to use a general language corpora and test sets for their experiments, several datasets are available, most of them in the English language. One of these freely available testing datasets for English is the Reuters corpus Reuters-21578 which contains documents from the Reuters newswire from 1987.¹ This corpus is mostly used for experiments on text categorization. The LSI website provided by Michael Berry and his group has published and linked several other corpora which can be useful for setting up LSA experiments.² The CISI corpus, the CRAN (Cranfield Collection) and the MED (Medlars collection) are even available as pre-formatted Term-Document-Matrices.

There are several other corpora available but most of them can only be used under specific license conditions or can only be accessed when asking for permission. One of these is the TASA corpus from the Touchstone Applied Science Associates (Landauer et al. 1998). The TASA corpus contains 10 million words of unmarked high-school level English text on Language, Arts, Health, Home, Economics, Industrial arts, Science, Social studies, and Business. It is often used in LSA experiments. If you use languages other than English, it is much harder to find appropriate general language corpora. A good starting point is the TRACTOR-service provided by

¹<http://kdd.ics.uci.edu/databases/reuters21578/reuters21578.html>

²<http://www.cs.utk.edu/~lsi/corpa.html>

University of Birmingham.³ Although a small access fee is charged and although most corpora are tagged there is a good chance that a usable corpus can be found here. The internet is another option for collecting sufficient written text material for performing an LSA. Especially the Wikipedia is a commonly used basis for experiments (Zesch et al. 2008).

Next to a general language training corpus, a second corpus is needed to add domain-specific concepts to the latent semantic space. This second corpus should contain sufficient material to represent the domain in which the experiments are taking place in. This material can be taken from textbooks, journal articles or even websites. The joined corpora are used to create a semantic space with reduced dimensions. The test documents are projected onto this semantic space. From here on one can compare concepts and documents by studying their vectors, and in particular the angles of the vectors to look for similarities.

When working with small corpora, often the generic “language” background corpus is omitted and only a domain-specific corpus is used. In this case the domain-specific corpus has to serve both purposes: to train the semantic structure of the target language used and to teach the particularities of the domain of interest. This approach only seems to work under certain restrictions and easily degrades effectively below the performance of the pure vector space model (not only with unfavorable parameter settings). In the next part of the chapter we introduce and discuss a number of software applications for conducting placement experiments with LSA.

12.3 Applications for Placement Support Experiments

In this chapter we will discuss examples of applications that can be used for a placement experiment. After a short review of existing applications a PHP implementation is introduced and a placement Web Service built in the framework of the TENCompetence project.

For a good general overview of existing applications that can be used for LSA we recommend the chapter by Quesada (2007) in the Handbook of Latent Semantic Analysis. The first software application for LSA stems from Bell Communication Research (Bellcore) which have also invented and developed the method. The company has changed to Telcordia and patented the method in the Unites States. The software consists of a set of UNIX applications that can be combined to do a full LSA based on command line expressions. An important LSA implementation stems from the group around Berry from the University of Tennessee. Their General Text Parser (GTP) was and still is the reference implementation for conducting LSA experiments (Giles et al. 2003). GTP supports the whole LSA process of constructing a TDM, applying an SVD and querying the latent semantic space. Another advantage of GTP is that it is available for several operating systems. But there are some drawbacks for using GTP as well. During the building of the matrix huge

³www.tractor.bham.ac.uk

temporary files are built and depending on the size of the corpus the user can run into memory problems. Another serious problem which cannot be solved by more powerful hardware is the incompatibility with 8-bit characters. This aspect makes it especially hard to use GTP for some European Languages which use diacritical characters. GTP is available upon request directly from Berry.⁴

Another application that we worked with is the MATLAB implementation for LSA (and several other techniques). The Text to Matrix Generator (TMG) was developed by Zeimpekis and Gallopoulos (2005) from the University of Piraeus. In their application the user can conduct a complete Latent Semantic Analysis including processing of texts, LSA space construction and querying. TMG has an easy to use user interface which offers a step-by-step procedure for a full analysis. A disadvantage of TMG is that the commercial product MATLAB is needed to run an analysis with it.

Another more recent implementation of LSA is the R package (Wild 2007). The R-framework offers a programming language and software environment for statistical computing. With the R package for LSA it is easy to set-up and run small LSA experiments with some lines of code. Although R does not have a GUI the code is easy to learn and well documented (for R in general and for the LSA package, too). One of the advantages of using the package are the several options to sanitize text within the R-environment and the easy import and export of data including statistical transformations. We will describe the R-implementation in the next part of the chapter more into detail.

12.3.1 The R Implementation

A typical LSA process using the R package looks similar to the one depicted in Fig. 12.1. First, a text matrix is constructed with `textmatrix()` from the input corpus. The text matrix can (but does not need to be) weighted using one of the various weighting schemes provided (see Wild et al. 2005 for more details). Then, the singular-value decomposition is executed over this text matrix and the resulting partial matrices are truncated and returned by `lsa()`. The number of dimension to keep can be set using various recommender routines (e.g., `dimcalc_kaiser()`). The resulting latent-semantic space can be converted back to text matrix format using `as.textmatrix()`.

In case that additional documents are to be folded into the existing latent-semantic space, again a new text matrix is constructed using `textmatrix()` re-using the vocabulary of the first. Again the resulting text matrix can be weighted (eventually re-using the global weights of the first text matrix). Using `fold_in()`, the resulting text matrix can be mapped into the existing latent-semantic space, thereby re-using the truncated left-sided and the diagonal partial matrices of the SVD. In this case, the result is directly a text matrix.

⁴<http://www.cs.utk.edu/~lsi/>

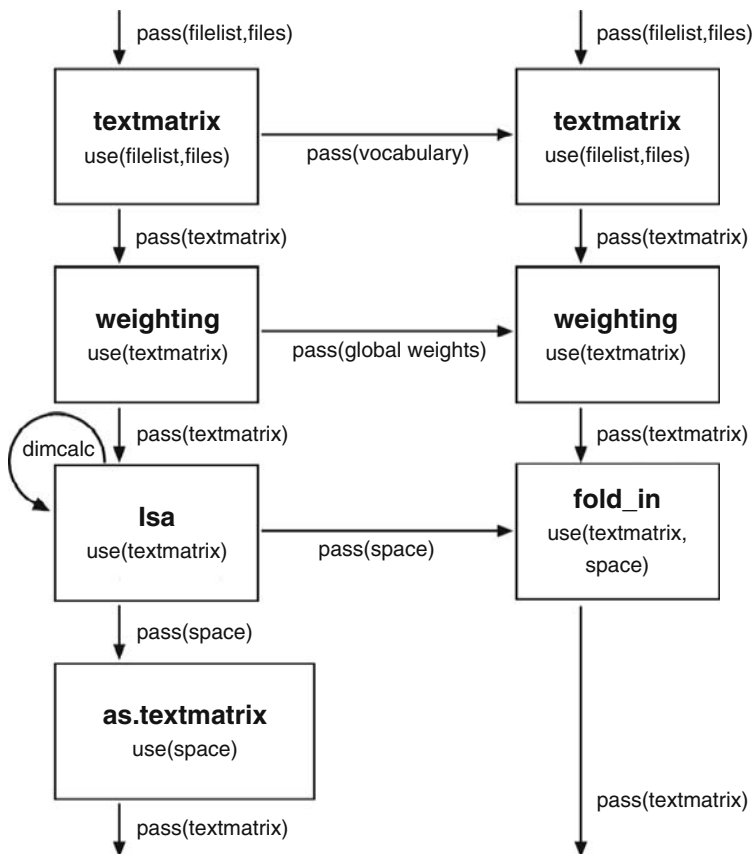


Fig. 12.1 Overall workflow

Looking more closely at the textmatrix routine, it can be seen that several text sanitizing and pre-processing steps are embedded in the text matrix generation routines: the routine included means to convert the terms to lower case, simple routines for stripping XML tags, automatic removal of punctuation marks and some other special characters, and trimming of multiple white spaces. Furthermore, stop words can be filtered (by providing stop word lists) or a controlled vocabulary can be deployed. Furthermore, frequency filters can be applied to delete terms below or above a certain frequency threshold (within a document or within the corpus) or outside a certain term-length range. Terms consisting purely of numbers can be removed automatically. Also all terms can be reduced to their word stems (using Porter’s snowball stemmer).

The following code example illustrates how LSA may be used to find the most relevant documents similar to a query document in a typical placement setting.

```

library("lsa")

#load training texts
trm = textmatrix ("trainingtexts/")
trm = lw_bintf(trm) * gw_idf(trm) # weighting
space = lsa(trm) # create LSA space
trm_red = as.textmatrix(space)

#fold-in query test document(s)
tem =textmatrix("test/", vocabulary=rownames(trm))
tem = lw_bintf(tem) * gw_idf(tem) # weighting
tem_red = fold_in(tem, space)

#find similar documents to test document E1.txt
cors = cor(tem_red[, "E1.txt"], trm_red)
which(cors>0.7)

```

Besides this implementation we have developed a PHP implementation for web applications. It is similar to a very recent implementation of LSA in Python.⁵ The Python implementation is a mixture between steps scripted in Python and using routines from a compiled math library. Until the time of writing this chapter we were not able to actually try out the Python implementation of LSA.

12.3.2 The PHP_Lsa Implementation

A prototype PHP implementation for LSA called PHP_Lsa was developed in the Cooper project (Ceri et al. 2007) where it was used as a tool to track down project documentation in company archives. This implementation (PHP_Lsa) is different from the existing applications in several ways:

- In contrast to Bellcore, Gtp, Matlab and R it is specifically made for use in web applications, although it can also be used in stand-alone/command-line mode. Besides the Python implementation and a web front-end using the R implementation, none of the implementation can be used easily in web applications.
- Matlab and R are more suitable for (manual) analysis purposes.

PHP_Lsa is more alike the recent Python implementation, although there are some significant differences:

⁵<http://www.joesniff.co.uk/projects/latent-semantic-analysis-in-python.html>

- PHP_Lsa is written as a compiled PHP extension where the Python implementation uses scripts to perform most of the actions needed to build the TxD matrix and perform queries and only uses compiled code for the SVD decomposition.
- The PHPLsa exposes compiled methods for all basic LSA steps from importing documents, through cleaning up, building the TxD matrix and decomposing and reducing it and performing queries. This significantly reduces the code complexity without removing the option to implement custom scripted steps in between.

Other differences are:

- It's possible to switch between a built-in and (two) external decomposition applications including win-gtp. No matter which engine is used, the output is always formatted as Harwell-Boeing matrix, making it easy to compare the various engines and to switch between them.
- PHP_Lsa also allows various locations for input documents like ftp, (local) disk and nntp (usenet). Other document sources can be implemented easily. It also has a wide range of textual cleaning functions built-in.
- PHP_Lsa can work directly with documents like Word, PowerPoint and Adobe's PDF. It also allows these type of documents as query.
- Because it saves all intermediate output automatically in an easy readable non compressed format if the matrices are small enough, PHP_Lsa is very suitable for research purposes. Because of the embedding in a scripting environment automation of certain actions like searching for optimal LSA settings is also easy to implement.

Of course there are also some downsides:

Like most LSA implementations PHP_Lsa is not suitable for very large corpora or documents. PHP_Lsa is only available on Windows and for specific PHP versions. The architecture consists of a number of layers. The bottom layer is the Sparse Matrix library that does perform all matrix operations including loading and saving in various formats. The native format for the Sparse matrices is Harwell Boeing (*.hb) (Duff et al. 1992). If the product of the rows and columns is small enough, the matrices are also automatically saved in a non-sparse format (*.ns) and Comma Separated Values format (*.csv) enabling further analysis.

On top of the Sparse Matrix Library the Lsa engine is built that performs the various actions necessary for Latent Semantic Analysis. These steps are:

1. importing documents
2. extraction of plain text from Microsoft Word, Microsoft PowerPoint and Adobe Acrobat documents
3. pre-processing the imported text
4. constructing the Term x Document matrix
5. performing the decomposition and transforming the output of the decomposition
6. reducing the rank of the Term x Document matrix

7. building a query vector
8. projecting a query vector onto the reduced Term x Document matrix
9. performing the query by calculating distances, dot products and cosines between the query vector and the reduced Term x Document space
10. allow retrieval of results

The last layer is the interface to the environment where the Lsa engine is used. In the case of the Cooper project this is PHP 5.2.0 or 5.2.4 for Windows XP running either from the command line or under an Apache 2.3.2 based web server.

The first step is to setup a PHP environment including a web server like Apache on a Microsoft Windows machine. The easiest way is to download EasyPHP2.0b1 from <http://www.easyphp.org>. This version contains PHP version 5.2.0. Due to lack of backwards compatibility between the extensions API in the various PHP versions PHP_Lsa is specifically compiled for PHP version 5.2.0 up to 5.2.4. Once the setup is completed the only things to do before starting to use the PHP_Lsa engine are:

- Copy the PHP_Lsa engine and optionally the external decomposition applications (wingtp.exe and/or svdlibc.exe, both which can be obtained from <http://lsa.colorado.edu> and <http://www.cs.utk.edu/~lsi>) into the php5\ext directory.

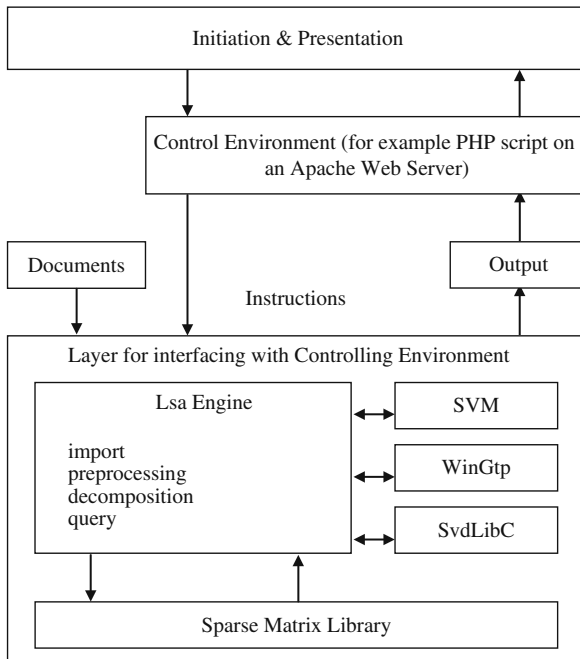


Fig. 12.2 The PHP_Lsa engine architecture

- Enable the PHP_Lsa extension in the PHP.ini file
- Verify PHP is running correctly and PHP_Lsa is enabled with `php_info()`. A small web page containing only `<?php phpinfo(); ?>` is normally used for this. It should show a section like:

Latent Semantic Analysis for PHP	
Lsa Version	Version 1.0 (Build 73)
Lsa Engine	c:\inetpub\php5\ext\php_lsa.dll
Lsa SVD Engine	c:\inetpub\php5\ext\svd.exe
Lsa Temporary BaseDir	C:\WINDOWS\TEMP\
Lsa Class	php_lsa

After this PHP and the PHP_Lsa engine are ready to be used.

Note: It is also possible to install PHP and PHP_Lsa under Microsoft IIS web server. PHP_Lsa was tested under Windows XP running IIS 5.0 with PHP version 5.2.4 installed. On top of this PHP implementation of LSA we have build an exemplary placement web service, which we will introduce shortly.

12.3.3 The Placement Web Service Prototype

In the context of the TENCompetence project a webservice was developed based on the PHP_Lsa application. This webservice is embedded in the so called “personalization pipeline” which has to guarantee that the optimum way through a learning network is recommended to the learners. The placement service is located in the bottom up cluster of the personalization pipeline. In this pipeline the positioning service analyzes portfolio data and projects them into a latent semantic space build from general language data, domain data and the content related to learning activities in Learning Networks. The result of the placement service is used by the personalization pipeline to recommend a set of learning activities which fit to the learner’s preferences and prior knowledge. In this case learning activities with a very high cosine similarity are omitted for this learner.

From a technological perspective the web service has been tested regarding reliability and performance. The division between the construction of the latent semantic space and the query process has improved performance for the service a lot. But the query process takes still around 30 s for a test dataset of 800 documents until the service is able to present results (total size 7 MB). This depends to a large extent on the corpus and query size.

In the Cooper project technical validation of the PHP-Lsa was conducted that will not be reported here. Work to validate the LSA-based placement service was done using a data set from introductory psychology at the Open University of the Netherlands. A part of that validation work is reported in Chap. 13.

Personalization interface -> Placement service

Calculated similarity between portfolio documents and learning network content

Name	Method	Description	Input (Parameter)	Output
getPositionValues	Get	Return a list of UoL annotated with cosine values	Iduser=xx	2-dimensional array of floats. Each UoL with its calculated cosine values.

Frequency	DATA fields	Format
When visualization tool has to be re-ordered or reorganized.	Iduser = Integer Learninggoal = Array (Strings)	-

12.4 Implementation

After the corpus has been prepared, the actual Latent Semantic Analysis takes place. For this purpose the prepared corpus needs to be processed with the chosen software application. It makes sense to think about a good folder structure for corpus data, temporary results, and final results. Since most experiments will need several runs before the optimal parameter settings have been found a good documentation system is needed in addition to keep the overview.

An first decision to take is whether weighting is applied. The raw frequencies in the term–document matrix can be transformed using local and global weighting. An example of a local weighting transformation is a binary transformation in which the frequencies are replaced by a binary value indicating whether a term is present or not. Another local weighting function replaces raw frequencies with their (natural) logarithm. Local weighting can be combined with global weighting that expresses the term occurrence in the corpus. Global weighting is applied to all cells corresponding to a term. Examples here are normalizing transformations, for example to normalize term vectors, and transformations such as inverse frequencies or entropy-measures. A combination of a local log transformation with a global inverse entropy weight is recommended by Quesada (2007) as yielding good results in retrieval and cognitive modeling .

One of the most important decisions in any LSA is the selection of the number of singular values (i.e., the number of dimensions) that will be used to reproduce the data. Dimension reduction itself is not a goal of LSA and it has become accepted practice that for corpora with the size of 5000 documents and above, at least 300 dimensions are used. Larger numbers are not uncommon. As Deerwester et al. (1990) already noted, it is core to LSA to not retain all singular values, because the “latent” factors only emerge in a model with lower dimensionality than the original. One method that might be used to decide on the number of singular values is

the Screen test . In this test one inspects the data to find the place where a sudden drop occurs in the size of the singular values. This point is used as a cut-off point, beyond which additional singular values are believed to add little more than error to the data.

A second approach is to normalize the document vectors to unitary length, and retain only the singular values that have a length greater than one. This corresponds to common practice in factor analysis of retaining only the eigenvalues larger than one. However, to routinely apply normalization to the data, would cause the risk to lose all information related to the length of documents.

The third approach, as proposed by Wiemer-Hastings and Graesser (2000) is to empirically determine the number of dimensions, for example, by selecting the number of singular values that optimizes a performance criterion, such as the correlation with ratings by human raters (Wiemer-Hastings 1999).

Kalz et al. (submitted) suggest to combine a performance criterion with a constraint on the amount of variance explained by the singular values. Their reasoning is based on the close connection between singular values and the variance they account for in the term-document matrix. We refer the reader to Chap. 13 for further details.

12.5 Conclusion

In this chapter we have introduced the basic tools and technologies needed to start experiments in placement. We have discussed several decisions that have to be taken and explained some reasoning from our recent experiments. Doing research with LSA means to simultaneously optimize the algorithm according to the application context, the language and the domain used in the experiment. Not many experiences and features of using LSA for educational purposes can be generalized. For the use of LSA for placement support ideally LSA can be trained to the decision pattern of experts which are responsible for APL procedures. We refer the reader Chap. 13 for further details

The main advantage of LSA is that it allows for fairly intelligent operations to be performed by putting in a minimum amount of effort. This is illustrated by its ability to use word co-occurrence data to move words and documents into a reduced dimensionality space where they can be more meaningfully compared to each other. This is fully automatic and does not require the use of meta-data, preliminary construed dictionaries, semantic networks, knowledge bases, grammatical syntactic analyzers, etc. Automation of processes can alleviate human workload considerably which is an additional advantage.

Because LSA uses a mathematical representation of the relations between words in a text and the semantic distance between texts it allows a rapid analysis of large numbers of documents. The strength of LSA is that it represents a corpus in a k -dimensional space and thereby reducing the complexity, which makes it possible to improve the estimates of pair wise similarities.

The weakness of LSA lies in the empirical determination of computational factors, the computational time that is needed to analyze big corpora, the directionality

of LSA and the application in contexts where the emphasis is on logic and (temporal) reasoning. Large matrices like Landauer et al. (1998) and others propagate may not always be needed but computation time may become a problem if the nature of material is such that it requires frequent re-computation of the underlying SVD. This is the case if there are large numbers of updates to the corpus.

References

- Ceri, S., Matera, M., Raffio, A., Spoelstra, H.: Flexible processes in project-centred learning. In *European Conference on Technology Enhanced Learning*, ed by Duval, E., Klamma, R., Wolpers, M. Lecture Notes in Computer Science, vol. 4753 (Springer, Berlin, Heidelberg, New York 2007) pp. 463–468
- Deerwester, S. et al.: Indexing by latent semantic analysis. *J. Am. Soc. Inf. Sci.* **41**, 391–407 (1990).
- Dessus, P.: Simulating student comprehension with latent semantic analysis to deliver course readings from the Web. *Cogn. Syst.* **6**, 227–237 (2004)
- Duff, I. S., Grimes, R. G., Lewis, J. G.: *User's Guide for the Harwell-Boeing Sparse Matrix Collection* (Release I). Technical Report RAL 92-086 (Rutherford Appleton Laboratory 1992)
- Foltz, P., Laham, D., Landauer, T.: Automated essay scoring: applications to educational technology. *Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications*. AACE, Chesapeake, VA, 1999, pp. 939–944
- Giles, J., Wo, L., Berry, M.: GTP (general text parser) software for text mining. In *Software for Text Mining, in Statistical Data Mining and Knowledge Discovery*, ed. by Bozdogan, H. (CRC Press, Boca Raton 2003) pp. 455–471
- Graesser, A. C., et al.: Using latent semantic analysis to evaluate the contributions of students in autotutor. *Interact. Learn. Environ.* **8**, 129–147 (2000)
- Kalz, M.: Positioning of learners in learning networks with content-analysis, metadata and ontologies. *Interact. Learn. Environ.* **15**, 191–200 (2007)
- Kalz, M., Van Bruggen, J., Rusman, E., Giesbers, B., Koper, R. (submitted). Latent Semantic Analysis of Small-Scale Corpora for Placement in Learning Networks. *Interactive Learning Environments*
- Laham, D., Bennett, W., Landauer, T. K.: An LSA-based software tool for matching jobs, people, and instruction. *Interact. Learn. Environ.* **8**, 171–185 (2000)
- Landauer, T. K., Foltz, P. W., Laham, D.: Introduction to latent semantic analysis. *Disc. Proc.* **25**, 259–284 (1998)
- Landauer, T.: LSA as a theory of meaning. In *Handbook of Latent Semantic Analysis*, ed. by Landauer, T. et al. (Erlbaum, Mahwah 2007) pp. 3–35.
- Quesada, J.: Creating your own LSA spaces. In *Handbook of Latent Semantic Analysis*, ed. by Landauer, T. et al. (Erlbaum, Mahwah 2007) pp. 71–85.
- Steinhart, D. J.: Summary street: an intelligent tutoring system for improving student writing through the use of latent semantic analysis, 2001. <http://lsa.colorado.edu/papers/daveDissertation.pdf>. Cited April 30 2008
- Van Rosmalen, P. et al.: Knowledge matchmaking in learning networks: alleviating the tutor load by mutually connecting learning network users. *Br. J. Educat. Technol.* **37**, 881–895 (2006)
- Wiemer-Hastings, P.: How latent is Latent Semantic Analysis? *Proceedings of the Sixteenth International Joint Congress on Artificial Intelligence*, San Francisco: Morgan Kaufmann, 932–937 (1999)
- Wiemer-Hastings, P., Graesser, A.: Select-a-Kibitzer: A computer tool that gives meaningful feedback on student compositions. *Interact. Learn. Environ.* **8**(2), 149–169 (2000)
- Wild, F.: R package for latent semantic analysis, 2007. Available at <http://cran.r-project.org/web/packages/lsa/index.html>

- Wild, F., Stahl, C., Stermsek, G., Neumann, G.: Parameters driving effectiveness of automated essay scoring with LSA. In *Proceedings of the 9th CAA Conference*. (2005), Loughborough, UK: Loughborough University
- Wolfe, M. B. W., et al.: Learning from text: matching readers and texts by latent semantic analysis, *Disc. Proc.* **25**, 309–336 (1998)
- Zeimpekis, D., Gallopoulos, E.: TMG: a MATLAB Toolbox for generating term-document matrices from text collections. In *Grouping Multidimensional Data: Recent Advances In Clustering*, ed. by Kogan, J., Nicholas, C., Teboulle, M. (Springer, Berlin Heidelberg, New York 2005) pp. 187–210
- Zesch, T., Müller, C., Gurevych, I.: Extracting lexical semantic knowledge from Wikipedia and Wiktionary. In *Proceedings of the Sixth International Language Resources and Evaluation. LREC'08*, Marrakech, Morocco, May 2008

Chapter 13

A Validation Scenario for a Placement Service in Learning Networks

Marco Kalz, Jan van Bruggen, Bas Giesbers, Ellen Rusman, Jannes Eshuis, and Wim Waterink



13.1 Introduction

In this chapter we describe a scenario for the validation of a placement service in Learning Networks. In Chap. 12 of this volume we described placement in Learning Networks as a case of Accreditation of Prior Learning (APL). We explained that placement in Learning Networks cannot assume the availability of data or metadata that allow for a direct or indirect coupling of data, such as competence descriptions, to the outcomes of learning activities. Even though such data may be available, their semantics are unknown, since there is no controlled vocabulary in Learning Networks. Given this state of affairs we have to operate under the strong assumption that similarity between material in the student portfolio and material in the target learning activities, typically course content, can be used as an indicator that the prior learning experiences of the student match the learning outcomes of the target

M. Kalz (✉)
 Centre for Learning Sciences and Technologies, University of Maastricht,
 Tongersestraat 53, 6200 MD Maastricht, The Netherlands
 e-mail:marco.kalz@ou.nl

learning activities. We elaborated the assumptions in terms of the assessment model of Pellegrino et al. (2001) and saw that this is not a trivial assumption.

We described the four phases of a typical APL procedure: (1) candidate profiling; (2) gathering and presenting the evidence, where evidence is collected about previous qualifications and experience that may support the claim that the candidate already possesses knowledge and skills that are outcomes of the learning program. This evidence is then presented in a portfolio; (3) assessing the evidence, where an assessor reviews the quality of the portfolio and (4) accreditation where verification or endorsement by the department responsible for awarding the credit or recognizing the positive outcomes of the assessment takes place. The second and third phases are considered the most critical: students do not know what evidence they could or should bring forward. Interpretation and assessment of the contents of the portfolio is problematic and time-consuming.

In Learning Networks we aim for a placement service that, at least in principle, can prepare a decision on exemptions for one or more target learning activities, comparable to step 3 in the APL process. It depends on the policies in the Learning Network and authority delegated by the provider of the learning activities whether or not the system can do the actual accreditation (phase 4). Before such a complex service can be offered, interim solutions seem more realistic. We can, for example consider a service that assists the new member of the Learning Network in the preparation of the portfolio by signalling relevant and irrelevant documents. Of course, such a service – assuming that it will reduce the content of the portfolio! – will also reduce the load of human assessors.

APL offers methods and techniques to identify prior learning experiences from formal and informal education. This procedure is especially important if a person crosses the boundaries between work and learning or between academic disciplines. Most of the methods for APL rely on experts who study the contents of the portfolios and decide which parts of educational programs could be exempted and which ones are best suited as starting point for the students. However, this way of analyzing prior learning is a very time-consuming and expensive method (Skinner 2005). Therefore, we are looking for methods and applications that may offer technological support for this process.

Although the validation scenario described in the remainder of this chapter is applicable to placement services that are based on different assumptions and techniques, we have illustrated and elaborated the scenario using the particular placement service that we are working on. The work we present is work in progress, rather than stable research results.

13.2 Validating a Placement Service in Learning Networks

The validation scenario we describe here is meant to test the *viability* of a placement service that helps to reduce the work load of assessors by filtering or annotating the contents of a portfolio. In the description of the scenario we keep the technical aspects of the LSA applications used to the essential minimum, although some of the parameters used have such an impact on the results, that we have to discuss them.

The validation scenario we are using demands that the placement service meets the following criteria:

1. *Scalability*: the service offers a roadmap to placement service that can use large corpora. This is a criterion clearly related to our goals of growing into an LSA application using large underlying corpora. Currently, there still is no large Dutch corpus readily available on which one could build a domain corpus that would allow training an LSA-based service (Louwerse and Van Pier 2007). LSA is often based on large scale training corpora for a domain of 10,000 and more documents. Opposed to that, in learning networks we will have to deal with corpora that are considerably smaller. One may question whether the size of the corpus is a problem, since domains such as Veterinary Medicine or Psychology are potentially very large. However, the learning network will often contain specific subsets of such domains, such as “Diagnostics” or “Psychology on a Bachelor level”. Moreover, the content encountered in learning networks is often limited to learning material, rather than having complete sets of reference works and background documents included.
2. *Sensitivity*: the placement service is capable to identify material dealing with the same topic, whilst discriminating material that is dealing with different topics. In the next sections we describe the preparation of a test corpus on which we tested the sensitivity of an LSA-based placement service. As we learned, the hard way, this is a crucial step in the strategy.
3. *Reliability and validity*. A placement service needs to be *reliable*, that is treating similar cases in an equal manner. In practice this means that the scores or ratings assigned to a document by the placement service need to match, within reasonable boundaries, those of a human observer. There is an upper limit here; it is highly unlikely that two human assessors would always rate student’s input exactly the same; their inter-rater reliability will be less than 1.00. In practice, one may be quite satisfied with an observed inter-rater reliability of 0.75 and higher. The observed inter-rater reliability is the maximum that one may expect from the ratings of a placement service. *The validity* of a placement service is the extent to which the service actually measures what it is supposed to measure. As the reader will have noticed, this a crucial criterion because the placement service is based on the assumption that content can be used as a proxy to learning outcomes. Measuring similarities between texts is not an issue, various standard solutions to achieve that are available, but whether these similarities correspond with similar learning outcomes is the validity question. We describe in some detail the steps we took in a field study to determine the reliability and validity of our placement service. The results reported are those of work in progress.
4. *Fit in APL procedures*: the placement service needs to fit in regular APL procedures. This implies that, first we need to establish how the service can collaborate with students and assessors in such a way that the work load, in particular of assessors, is reduced. This has implications beyond considerations of reliability: the service needs to be optimized so as to avoid errors that are costly in terms of assessor time. Second, the placement service should use the results of

the APL process to improve its performance. Thus, ratings of assessors of a particular case can be fed back into the service by using the material of the case as a positive or negative example. In our LSA-based placement service this type of case-based reasoning is achieved by adding these examples as new positive or negative standards to which student data can be compared.

In the next sections we present and discuss our work to meet the requirements set forth in this validation strategy.

13.2.1 *Testing the Sensitivity of the Service*

For a placement service that uses similarity of textual materials as a proxy to similarity of learning outcomes it is necessary that it is able to identify material that deals with similar topics whilst discriminating material that deals with different topics. To that end we construed a test corpus with the type of learning material that we expect in Learning Networks: material that shares a general theme, but represents various topics and has substantial overlap between the documents (as if originated from different providers). As a theme we selected “apes and monkeys” and we created a corpus by collecting documents from the Dutch version of Microsoft Encarta and a Web search. These documents were split by the authors according to the topics treated in the text. The preparation and cleaning of the corpus was a tedious process of which more details can be found in the textbox.

Preparing and cleaning a test corpus

We created a test corpus of materials with characteristics corresponding to what we expect from learning materials made available in Learning Networks. We created a corpus in a domain where the material would be primarily *factual*, so as to facilitate text comparison. Second, we decided to use sources that would be mainly *textual* in order to minimize pre-processing. Third, we decided, more or less random, to build a corpus around ‘apes’ (the Dutch term ‘aap’ refers to apes as well as monkeys; we will use ‘apes’ in this broad sense). In order to keep the size of the corpus limited we choose to create the corpus in Dutch.

The corpus is based on a collection of documents obtained from the Dutch version of Encarta encyclopaedia and from a Web search (via Google Search) of documents in Dutch on ‘apes’ and ‘monkeys’. This resulted in a series of documents, including a number of overview articles, in which several species of apes were described.

Cleaning the corpus

Although it would lead to (slightly) better results, we decided to create a ‘clean’ corpus and all analyses were done after pre-processing the documents

in the corpus. We used simple tools such as TextStat (Hüning 2005) and TextPad to do most of the cleaning of the corpus, that is remove duplicate documents, repetitive misspellings, diacritical signs and repair split tokens.

Running statistics such as word frequencies revealed a number of spelling errors in the corpus. These errors were corrected only when a document contained a number of the same errors, such as systematic misspelling as, e.g. 'oran utan'. For further analyses the text of documents in the corpus needed some additional preparation, in particular removal of diacritical.

Splitting the documents in the corpus

The test corpus was created with the purpose to test the sensitivity of the model, i.e. discriminate between texts that treat different topics, as well as recognize texts that treat similar topics, here different species of apes and monkeys. Many documents in the corpus presented information about families of apes or monkeys. These documents were split in fragments dealing with a particular species only. Obviously, text fragments of a more general nature, such as texts dealing with general characteristics of the big apes, would remain. These documents were kept in the corpus.

The splitting was done manually by two researchers who tried to identify segments dealing with one topic. Wherever possible, splitting was based on text structures, such as headers and paragraphs in the text. If a split would result in a document with less than 20 words, no splitting was done. All splitting was discussed and agreed between the authors. Splitting the documents in this way created new documents of about a paragraph in length.

The final mean document size was 311 words. The final corpus contained 287 documents and had a size of 611,216 bytes. TextStat (Hüning 2005) reported 10,736 different tokens of which 4376 occurred in at least two documents.

Towards automatic splitting of documents

Splitting the documents into fragments of about a paragraph size has become more or less common practice in the application of LSA. Manual splitting as was done here is obviously not a viable solution in the context of learning networks. An alternative – not tested in the creation of this corpus – is to detect topic change by correlating subsequent sentences and treat a drop in correlation as an indication for topic change.

Filtering the corpus

The test corpus used is rather small and its data – expressed in a Term-Document matrix with term frequencies in the cells – can be expected to contain a substantial amount of error. One way of reducing the error in the matrix is to exclude terms from the Term-Document matrix that do not contribute to the systematic variance. The primary candidates are terms that have

a high frequency in the corpus. “Stopping” these terms can reduce error from the data before analysis.

However, indiscriminate filtering of high frequency terms in a small corpus may have the side-effect of removing domain specific terms that are frequently used in the particular corpus. Obviously, terms like “monkey”, “ape”, but also “food” occur with a higher frequency in the corpus than they do in regular Dutch and these domain-specific terms we prefer to maintain.

In order to prevent domain specific terms being filtered away solely on the basis of their frequencies, we compared the rank orders of the terms within the corpus to those within a more general corpus, based on several volumes of a Dutch newspaper (Bouma and Klein 2001). Terms with the highest ranking in the corpus were selected for stopping, unless the difference with the rankings in the general corpus exceeded 50 ranks (an arbitrarily chosen number). Thus domain specific terms that are “overrepresented” in the corpus will be retained. Thus, the term “aap” (ape or monkey) that would be removed from the corpus if we only considered general word frequencies, is now being retained using this filtering method. For analysis purposes a number of variants were made that would filter 0%, 30%, 40% or 50% percent of the occurrences of terms in the corpus.

Stemming the corpus

Stemming is the process in which a token is reduced to its stem, ‘Hypotheses’, ‘hypothesis’ and even ‘hypothesize’ may be reduced to a single stem. No stemming was applied to the corpus for the practical reason that we had no access to a quality stemmer. Languages like Dutch and German seem good candidates for stemming, because they tend to concatenate words to create new tokens. For example, whereas in English two tokens are used, e.g. ‘mountain gorilla’ for a particular species, Dutch and German concatenate the tokens, in this case two stems, to create a single new one: “berggorilla”; in a similar way the concept “school achievements” is represented in Dutch by the single token “schoolprestaties”. This concatenation process has the interesting side effect that the tokens in the corpus are more context specific than would be the case in an English counterpart.

We tested the sensitivity of the service by comparing the mean correlation of documents *within* a homogeneous set of documents (all dealing with a particular species) to the mean correlation of documents *outside* this set. One homogeneous set would for instance hold a collection of documents dealing with gorillas; another would consist of documents about chimps et cetera. We then tried to determine a point of optimum performance, where maximum correlations between documents *within* the homogeneous set coincide with minimum correlations between documents *outside* the subset.

In the analyses we varied strategies of filtering the data, as well as the number of singular values used to reproduce the data. All analyses were done with the General Text Parser (GTP) program (Giles et al. 2001). The singular value decomposition used *raw* frequencies. Other types of frequency measures, such as inverse frequencies or entropy based measures have the drawback that they will filter away domain-specific terms with relative high frequency in a small corpus like this. Document vectors were not normalized, but most of the documents in the corpus are of similar length anyway. Cosine values were used as association measures. The summary results of one contrast under one filtering strategy are plotted in Fig. 13.1.

Figure 13.1 plots correlations and explained variance (for an explanation see the textbox) obtained under various numbers of singular values (x-axis). The two descending curves in the figure represent the mean correlations within (squares) and outside (dots) the homogenous set. If we calculate the differences between the two means we have a measure for discrimination. Taking this as our performance criterion we could, following Wiemer-Hastings and Graesser (2000), select the number of singular values that optimizes discrimination.

However, as the plot shows, the discrimination is best when only a few singular values are used. This is much an artefact of the analysis method used and in order to compensate for that we introduce boundaries of minimum and maximum amounts of variance explained. The rationale for that is rather simple: the more singular values we use, the better the original data are reproduced, but that includes the error contained in the data as well, and that will lead to diminishing correlations (as the plot shows as well). In our validation scenario we use the amount of explained variance to determine the range of singular values (see the x-axis) that is used to determine the optimum performance. Typically we apply a range of 75–90% of the




Fig. 13.1 Mean correlations and explained variance

variance explained. The rising line in the plot represents the proportion of variance explained and we can now see that if we use only a few singular values this explains about 50% of the variance. Once we combine performance (that is discrimination) with a selection of the number of singular values that will account for 80–90% of the variance the range of singular values to be considered is 35–45.

Singular values and variance accounted for

There is a close connection between singular values and the variance they account for in the Term-Document matrix analyzed by LSA. The sum of the squared singular values is equal to the sum of squared cell frequencies (Fröbenius norm). Since the Term-Document matrices are sparse, that is, contain lots of zero-filled cells, the mean of the cell frequencies is close to zero. But, since the mean is close to zero and N the number of cells, is large, the variance is determined mainly by the sum of squares of the cell frequencies. We can therefore estimate the proportion of variance accounted for by using the squared singular values only.

Thus, one may select a minimum and a maximum number of singular values that corresponds with a bandwidth of variance accounted for. In an ideal world, one would retain the number of singular values that corresponded to the reliability of the data. This would ensure that the maximum systematic variance is retained, whilst the maximum of error is being removed. Not knowing these reliabilities we can only make an attempt at reducing error in the data. We do so by filtering the data before the analysis (stopping) and by limiting the number of singular values that are used to reproduce the data after the analysis.

13.2.1.1 Conclusion

For this part of the validation scenario we conclude that the approach is *sensitive*, in that it is able to recognize documents with similar topics and discriminate between documents with different topics. This sensitivity could, however, only be demonstrated when we filtered the corpus. Essentially, filtering reduces error and we presented a heuristic that does so while retaining high-frequency domain-specific terms. Essential in the approach is the optimization of a performance criterion within boundaries of explained variance, that correspond to expectations about the reliability in the data.

There are some limitations and caveats to mention. First, we presented one sample of the analysis. In real validation the results should be replicated using several sets of documents. Second, we worked with a corpus that was sanitized by correcting systematic misspellings and other errors. In reality, the data will be less clean and there will be no opportunity to prepare the corpus in this way. Third, this validation step assumes document splitting that is based on topical content. For a real implementation the splitting should be done automatic.

13.2.2 *Testing the Reliability and Validity of a Placement Service*

The next step in the validation scenario is to demonstrate that the placement service has sufficient *reliability* and *validity*. Reliability can be understood here as a quality to “treat similar cases in the same way”. Whereas it is hard for human assessors to apply assessment criteria consistently over lengthy periods of time, for computers this no problem whatsoever. ‘Similar’ however is not the same as ‘equal’ and for that reason it is still necessary to compare computer ratings to that of human experts. Validity is the quality to measure what one intended to measure. For the placement service the determination of validity is crucial, because our approach is based on the assumption that we can use similarity of documents as a proxy to similarity of learning outcomes.

In this section we describe our method to test the reliability and validity of the placement service. The *method* is of prime importance here; the data – limited as they are – are presented to show how the approach works (and the lessons we learned from applying it).

We collected data among students and staff engaged in an introductory psychology course at the Open University of the Netherlands. The course uses a course book of 18 chapters each of them dedicated to a subtopic in psychology. The department of Psychology is interested in adapting the introductory course to the needs of its students. Methods of adaptation being considered are flexible sequencing, in which students are recommended to take the chapters of the course book in a different order, as well as accreditation of prior learning, in which case students can be exempted for particular chapters.

The introductory psychology course is offered in an online environment. Before they could enter the actual course students had to read an introduction about the content of the units of the course. They then were asked to fill out a questionnaire on prior knowledge for the course and the individual chapters of the course book. To substantiate any claims regarding prior knowledge students were invited to submit relevant material that they had produced in prior education or working environments. From 200 students that enrolled in the course 15% submitted material as evidence for prior learning. The documents submitted included work reports, theses and technical reports.

The documents were analyzed in two steps: first, the similarity between the student documents and the individual chapters of the course book was determined (a total of 234 ratings per observer); second, for each student document it was decided whether it contained sufficient evidence to warrant an exemption for one or more chapters in the course book. Two domain experts who were familiar with the course and the course book independently rated the similarity between each of the student documents and each of the 18 chapters in the course book. They also indicated whether or not a document entitled the student to an exemption for one or more chapters.

The observers rated similarity on 5 point Likert-scale. The raters showed an overall percentage of agreement of 87.6%. Inter-rater reliability measured through the coefficient of concordance showed a value for Kendall’s W of 0.59. Cronbach’s

alpha as a measure for the intra-class correlation is 0.82. While the raters showed an overall percentage of agreement of 87.6% the LSA results showed an overall percentage of agreement with the two raters of 59.6% and an inter-rater reliability measured through the coefficient of concordance value for Kendall's W of 0.59.

Corpus creation and analysis

In this study we also moved towards larger corpora. As a base corpus we used the PAROLE distributable corpus from the Institute of Dutch Lexicology. Parole distributable is a collection of modern Dutch texts mostly originating from newspaper or magazine articles with about 20 million tokens. Overall, the corpus consists of 4727 documents, including 800 selected documents representing the psychology domain and 18 documents containing the content of the introductory psychology course.

All analyses were done using the Text to Matrix Generator (TMG) which is a Matlab implementation of Latent Semantic Analysis and other techniques (Zeimpekis and Gallopoulos 2006) (See Chap. 4.4). Most documents were cut into pieces of 5 kb, only the target documents remained in their size. We tested several models to estimate the ideal number of factors which we found in the literature: 10%, the magic 300–400 factors, the estimation of the ideal numbers of factors based on the heuristic explained earlier in this chapter. Overall we did 28 analyses with different settings. In addition we did several runs with different stop word lists. We used a stop word list that reduced the corpus for 30% (20 stop words) and a stop word list that reduced the corpus for 50% (130 stop words). We applied normalization to the documents and we experimented with different weighting settings and local and global frequencies.

The limited reliability found in these data means that empirical verification of the validity of the placement service grinds to a halt here: validity can only be established if there is sufficient reliability. Closer inspection of the data revealed several problems: the raters seemed to agree most on student documents that had no similarity. As Table 13.1 shows, the majority of the cases ($n = 194$) is considered by both raters as having no similarity whatsoever to any of the chapters.

In the other cases the raters show less agreement with rater 2 in general scoring somewhat higher. Concentrating on the cases that showed similarities, according to at least one rater, we defined three cases: a perfect match of the ratings; a near match (1 point difference) and a non-match (2 points or more difference). Table 13.2 reports these (non-)matches.

The reliability of the human raters is at a satisfactory level. More important is the relation between similarity ratings and exemptions. In 85% of the cases where the raters rated a high similarity between a student document and a course book chapters, they also indicated that an exemption should be granted. This finding supports

Table 13.1 Similarity ratings of the experts

				Rater 2		
		1	2	3	4	5
Rater 1	1	194	9	9	3	0
	2	2	6	24	3	6
	3	3	0	3	0	3
	4	0	0	3	9	3

Table 13.2 Perfect, near and non-matching ratings

		Frequency	Percentage
Valid	Perfect	10	25.0
	Near	21	52.5
	Non-match	9	22.5
	Total	40	100.0

the basic assumption underlying the placement service described here, namely that similarity in materials can be used as a proxy for similarity in learning outcome.

So the final step left in testing the validity is to demonstrate that the placement service will rate the similarity of the student documents in a way that is comparable to human experts. Unfortunately, initial analyses showed that the correlation between LSA similarity scores and those of the experts were too low to use as base for a recommendation regarding exemptions.

There are several reasons why this result could occur. In the first place the data is distributed very uneven. The majority of the student documents show no similarity with the target material whatsoever, resulting in a highly skewed data distribution. Although this might be a realistic case for placement practices, it also means that a statistical model on which a decision (accept/reject the material as evidence) is based will be biased towards a general reject, which is a safe choice for almost every case. Additional work is needed to take into account these statistical properties of the data.

However, more important than statistical considerations is the crucial validity aspect, namely that the LSA seems to assess similarity in a different way than the experts do. There is some evidence that the similarity rating of the LSA depends more on the general Dutch language used, than on the domain specific elements in the material. The results of the analyses based on the psychology corpus alone, are nearly identical to the results obtained when the Parole corpus was used as well. What may have added to this effect is that long texts are being compared: student essays versus complete chapters. Salient characteristics may have disappeared under a large amount of general language usage.

Further inspection and analyses of the corpora are clearly needed. First, we need to test whether the semantic spaces of the Psychology corpus and the Parole corpus are almost identical indeed. If that is the case, then the corpus has to be modified in order to increase its domain specific part. Second we plan to experiment with smaller segments of source and target documents to check whether that will bring out the salient characteristics. Before and during all this we have to calibrate the service on the characteristics that the expert use to rate the texts in order to finish testing the reliability and validity of the placement service.

As the reader will have noticed these results indicate that for the domain of Introductory Psychology the problem may have to be traced back to a lack of sensitivity. One of the lessons learned here is to not take for granted that the service will be able to discriminate materials in the domain.

13.3 Fit in APL Procedures

In the preceding section we concentrated on the question how we could test the reliability and validity of a placement service in Learning Networks. It became clear that additional work is needed before the service can pass the test of this step. Yet, that additional work will also be guided by taking into account how a service could be incorporated in an APL process. We want to construe a service that helps to minimize the work load of human experts and that can learn from the decisions of the assessors who inspected portfolios.

The place in the APL process is that of a system that advises during the stage where a portfolio of evidence is put together on the relevance of documents submitted. As a preparation of the stage in which the portfolio is assessed the service annotates the documents as to their relevance for possible exemptions. To avoid introducing workload the system should be fine-tuned to avoid flagging documents as relevant where they are not.

Unfortunately, in the data of the Psychology course there were so many irrelevant documents that any statistical model would classify all documents as “irrelevant”. To at least partially overcome this problem we intend to create a feedback loop in which the system can also compare documents to documents that were previously classified by human raters. Thus documents that were accepted as evidence leading to an exemption, as well as documents that were rejected as evidence can be fed back to the system as standards to which new documents can be compared. The role of such a system would be left unchanged: it can scrutinize a portfolio and signal which documents bear no similarity to target material or prior positive cases. This would then result in a recommendation to remove the document from the portfolio. For assessors the service would annotate the documents in the portfolio with similar content and links to portfolio’s with relevant (positive or negative) documents that were assessed earlier.

13.4 Conclusion

In this chapter we presented a validation scenario for a placement service that is based on the assumption that similarity of learning material can be used as a proxy to similarity of learning outcomes. We emphasize that such a placement service reflects a worst-case scenario where no usable description whatsoever of the contents and the learning outcomes are available, or – as in the case of Learning Networks – cannot be assumed.

We described how such a placement service can be validated following a number of steps that we tried to illustrate with data and descriptions to give the reader a flavour of the work involved. As stated before, this chapter presented an overview of work in progress and we can now present some conclusions and lessons learned from this work.

First, we found some evidence that the assumption underlying the content-based placement service is valid: high similarity ratings coincide with exemptions granted. Second, we developed a number of heuristics to deal with small corpora, in particular dealing with filtering terms and determining the number of dimensions to retain. Third, we learned the hard way that running a complete cycle for the validation scenario is necessary. Sufficient sensitivity that corresponds to classifications made by human raters cannot be assumed. Finally, although we conceived the service as assisting and not replacing actors, it has become evident that more attention should be given to automatic corpus creation. Of particular interest here, is the automatic segmentation of documents based on semantic principles, rather than on an arbitrary maximum size. A way forward here is shown by Gounon and Lemaire (2002) who used LSA to identify “topical paragraphs” in a text. It remains to be seen whether this will also handle the problem of rating similarity on general language characteristics. Beside adding more domain specific material and topical segmentation we are considering feeding the system with (several copies of) metadata and related material, such as keywords, index terms, summaries, and learning objectives.

References

- Bouma, G., Klein, E.H.: *Volkskrant database and practicum* [Data file] (2001). Retrieved May 12, 2005 from <http://hagen.let.rug.nl/~hklein/Volkskrant/Volkskrant/> http://hagen.let.rug.nl/~hklein/pract_prainbow2.html
- Giles, J., Wo, L., Berry, M.: GTP (general text parser) software for text mining. In *Software for Text Mining, in Statistical Data Mining and Knowledge Discovery*, ed. by Bozdogan, H. (CRC Press, Boca Raton 2003) pp. 455–471
- Gounon P., Lemaire B.: Semantic comparison of texts for learning environments. In *Advances in Artificial Intelligence – IBERAMIA*. ed. by Garijo, F. J., Riquelme Santos, J. C., Toro, M. Lecture Notes in Computer Science, vol. 2527 (Springer, Berlin, Heidelberg, New York 2002) pp. 724–733
- Hüning, M.: *TextStat Simple Text Analysis Tool*. (Dutch Linguistics, Free University of Berlin, Berlin 2005). <http://www.niederlandistik.fu-berlin.de/textstat/software-en.html> Cited May 12, 2005

- Louwerse, M., Van Peer, W.: Waar het over gaat in cijfers: LSA als kwantitatieve benadering in tekst- en literatuurwetenschap [Where it's about in numbers: LSA as a quantitative approach in text and literature studies]. *Tijdschr. Ned. Taal Let.* **123**, 22–36 (2007).
- Pellegrino, J. W., Chudowsky, N., Glaser, R. (eds): *Knowing What Students Know: The Science and Design of Educational Assessment* (National Academy Press, Washington, DC 2001)
- Skinner, K.: *Continuing Professional Development for the Social Services Workforce in Scotland* (Scottish Institute for Excellence in Social Work Education, Dundee 2005)
- Wiemer-Hastings, P., Graesser, A. C.: Select-a-Kibitzer: A computer tool that gives meaningful feedback on student compositions. *Interact. Learn. Environ.* **8**, 149–169 (2000)
- Zeimpekis, D., Gallopoulos, E.: TMG: a MATLAB toolbox for generating term document matrices from text collections. In *Grouping Multidimensional Data: Recent Advances in Clustering*, ed. by Kogan, J., Nicholas, C., Tebouille, M. (Springer, Berlin 2006) pp. 187–210.

supporting continuous professional development with mobile and ubiquitous access to information, communication, and networking facilities. The chapter analyses best practices and gives the reader a reference model for designing mobile learning network services with 5 dimensions on the type of content, the information flow of the involved content entities, the context parameters taken into account, the pedagogical models involved, and the range of purposes of the services. Based on this reference scheme 5 best practice applications are presented and concrete examples are given to the user for applying mobile social software in learning networks.

Chapter 16 highlights a different shade of mobile access to learning networks, i.e. the question how authenticity in learning support can be implemented based on mobile user generated content or the mediated connection of real world situations and more formal learning settings. The chapter presents different examples ranging from enriching classical classroom situation to extending work or home environments with new media enhancements to ensure high quality learning in informal contexts and understand the impact of authentic learning support on intrinsic motivation, interest, engagement, immersion, or sustainability. The chapter gives successful best practices for different patterns enabling the support of distributed learning networks with the integration existing facilities (museums, outdoor parks, production facilities, expert workplaces, specialized service centers, and others).

Building on the previous argumentation Chap. 17 deepens a special aspect of mobile access to learning networks, which is considered with how to achieve reflection in action and social awareness with using peripheral (contextual) information. Therefore the chapter gives a sound introduction to the underlying concepts of awareness and reflection for learning, an analysis of peripheral information and interaction footprints, and contextual dimensions in relation to concepts of communities of practice. Thereafter the chapter presents team.sPod, a system designed, implemented and evaluated, which takes into account the complex relationships and aims at supporting reflection in action for continuous professional development.

In general the section highlights several aspects of contextualized learning network services and gives the reader best practices, a reference model, and a theoretical analysis of underlying principles and success factors for stimulating professional development across contexts.

design, and empirical research on technology enhanced learning. When classifying best practices of adaptive instruction different parameters of the instructional process have been identified which are adapted to the learner, as: sequence and size of task difficulty, time of feedback, pace of learning speed, reinforcement plan and others these are often referred to the *adaptation target*. Furthermore Aptitude Treatment Interaction studies explored the effect of adapting instructional parameters to different characteristics of the learner (Tennyson and Christensen 1988) as task performance, personality characteristics, or cognitive abilities, this is information is referred to as *adaptation mean*. Intelligent Tutoring Systems (ITS) discussed topics of acquiring information about the learning process, building cognitively adequate learner models, inference information about learners, and develop effective recommendation and guidance strategies for personalized learning paths. In ITS basically expert models of the problem domain have been build to give individualized feedback and support the development of problem solving competence in the target domains (Anderson et al. 1989; Carbonell 1970; Clancey 1987). User modelling (UM) research has focused on acquiring and representing information about the user/learner and also how to use this information for adaptive systems.

From early 1990s educational adaptive hypermedia systems (Brusilovsky 1996) have been using mostly simpler models of learner knowledge and preferences to adapt the presentation and selection of hypermedia content as: annotation of hyperlinks, sequencing of learning contents, or content recommendation (Brusilovsky 1997; Specht and Kobsa 1999). Early works taking into account social information for giving instructional guidance also came from the field of adaptive hypermedia as for example using information about the usage of learning content by other members of a learning community or peer group members (Brusilovsky 1996).

In the last 10 years the technology available and used by learners has dramatically changed. In this section we would like to introduce a shift from personalization towards contextualization in technology supported learning networks. A shift from taking into account learner knowledge, preferences, goals, cognitive abilities or task context of the learner towards the social and environmental context of learning has and is happening. For example in the area of adaptive hypermedia the knowledge sea project takes principles of adapting coloured navigation maps using information about the behaviour of peers and learner groups (Farzan and Brusilovsky 2005). More and more instructional systems have been designed and implemented in the last years that take into account other contextual parameters as the social environment, the location of the learner, or peripheral information not directly related to the subject matter but the context of learning. From the background of context-aware systems the relation between personalization and contextualization and new methods for building contextualized user models have recently been discussed in several articles (Zimmermann et al. 2005b). Beside new possibilities of understanding the context of the learner for adapting the learning environment also new possibilities of accessing information and documenting learning experiences via mobile technology have become essential (see Sect. 14.4). Learning content can be accessed from mobile and ubiquitous devices anyplace, anytime, and anywhere, but even

more not only the user's location, working context, and performance can be used to filter information but also important aspects of the environment and the available infrastructure can be used to enable interaction and contextualized learning experiences. Learning activities can be performed in real life situations and be connected to remote experts (Chap. 16), content infrastructures for social annotations (Chap. 15), personal reflection, awareness (Chap. 17), and indicator functionality enabling a distributed and at the same time embedded learning experience.

Adapting instruction to the learner is an intensively discussed topic but with a completely different background considering also recent developments in web 2.0 technologies. Learners have become mobile prosumers of information and use information about their social networks and peer activities from mobile devices or ubiquitous access to information. Recent studies about how a new generation of learners uses mobile and social information technology indicate a shift not only in usage, but also in the meaning of these technologies for learning (Green and Hannon 2007). The notion of adaptivity changes in the sense that the perspective has to be broadened from a computer system monitoring a user/learner and adapting to her behaviour towards learners that create their own flexible and dynamic content spaces where information needs, information provision, learning support, and consumption of support services are inherently interwoven in learning networks. In this chapter we argue that personalization is a core concept for delivering learning support but the classical view on personalization has to be extended towards contextualization.

In the following part we want firstly to describe a changed landscape for personalized instruction based on web 2.0 technologies, social software, learning networks, mobile software and context-aware and ubiquitous computing. In a second part we will describe a basic scheme for classifying adaptive instruction and highlight where the complementary potential of personalization and contextualization lies. In the following chapter the authors will describe different applications scenarios of mobile learning networks.

14.2 The Changing Landscape for Adaptive Learning

While 20 years ago the cold start problem, i.e. the access and availability of user information for building user models was a main question in user adaptive systems and adaptive learning environments, today user information is often available and even given freely from users as a natural part of self-marketing and presentation. User information for personalization is broadly available and even ubiquitous today; users leave tracks not only in the digital world but also in the real world where more and more sensor technology and the integration of technology in our daily life take place. Technology enhanced learning nowadays has more to tackle problems of merging user models, synchronizing and validating user information than to find solutions for acquiring basic user information (Heckmann 2006). Even more when this personal information is made available the underlying educational models like

e-portfolio see this information as a key element in building personal portfolios for presentation, networking, assessment and other functions and research their impact on learner variables as self-esteem and self-confidence (Hilzensauer 2007).

Also approaches for giving personalized recommendations have changed dramatically especially in the last 10 years with social software and social networks in the Web 2.0. Collaborative filtering in contrast to content-based recommender systems have become more and more popular and while we would still consider this as personalized support nevertheless the underlying algorithms for adaptation use completely different data and also have a different ratio when giving recommendations to a learner. Furthermore whilst in a first step the shift from user model and content based adaptation was done more and more sensor and context information is used in today's recommender systems. As we foresee in this chapter future recommender systems will take into account not only online click behaviour of users but there movements in physical space, tagging of places, recording of geo-annotations, or contextual blogging approaches to enrich real world objects with educational information (see Chaps. 15 and 16).

Along this line several essential trends can be identified along the path from personalized adaptive learning systems to contextualized learning networks including:

- *The Role of user information has changed:* Information about user activities come from web 2.0 social networks where users explicitly or implicitly leave traces and even connect systems collecting different information in e-portfolios and personal learning environments. Considering the value of different social software tools even meta-information about the learning process is given by learners and considered as a valuable resource for actual learning, as for example reflection about the own learning process in blogs. Recent studies have also shown the importance and effects of indicating context information to the learner for supporting self-organized and self-direct learning (see Chap. 17)
- *Services based MashUps build the new Personal Learning Environments:* New service-based architectures as personal learning environments describe a learning environment as an integration of loosely coupled services and personalization gives learners the possibility in adaptable environments to select preferred services and combine them to generate added value in mash-ups. Latest developments describe mashups not only as integration of different services but also stress the fact of distributed and mobile access to functionality and parts of a mashup or adapt access to this mashup controlled by the learning process. A recent overview is given in (Wild et al. 2008).
- *Pervasive learning technologies enable more authentic learning support:* Mobile learning support and embedding learning activities in authentic contexts becomes more and more important and has proven in the last years as an efficient method for supporting (see Chap. 16) sustainable and enriched learning experiences in a variety of educational sectors ranging from school field trips to life long learning support. Social software as a tool for creating content in context has established even in the corporate information landscape as a new mean of knowledge management and bottom up integration of the “wisdom of crowds”.

- *Ubiquitous and contextualized access to information and services enable the integration of informal and formal learning support:* Mobile and ubiquitous learning technology gives new possibilities to support distributed learning networks with activities ranging from notification for awareness, content creation in context and content exchange as also reflection in and about action (Schoen 1983). Furthermore the previous support for e-learning on a desktop computer de-contextualizing the learner from the actual community, practice environment, or working activities becomes dissolved when computers and through them learning services can be integrated in daily living and working environments. Through these technologies not only the information and services become available on-demand in a situation of practice but also the social networks and informal learning activities are integrated in a life long learning experience.

Considering these changes we will now first describe a model for describing adaptive systems (Specht 1998) and will highlight the impact and potential that the described changes have for adaptive, i.e. contextualized learning networks.

14.3 Classification of Adaptive Educational Methods

Basically, our understanding of contextualization is derived from an extension of the classical adaptive instruction approach on adaptation and personalization. In an analysis of a variety of adaptive educational systems we have developed a classification system for adaptive systems taking into account four dimensions of adaptive methods in educational systems (Specht 1998).

- *Adaptation Mean:* What information about the user is known, and what information can the system use for adaptation? For contextualized systems, this question can be extended to the user's environment and to new methods for integrating sensor data and making inferences about the context of use. For example, the system can use all of the sensors in a room in which a user is currently moving, or the system can use context information concerning related entities in another place. These new types of information allow for more valid and richer inferences in contextualized computing.
- *Adaptation Target:* What aspect of an adaptive system adapt to the given information about the user? For contextualized systems, this question can be extended to multimodal interaction and sensing. In classical user modelling, the user's environment is his/her PC desktop. The input channels for user tracking are mostly limited to the GUI interactions of the user. In ubiquitous computing, new channels for both user input and system output become available. Location can be seen as just one parameter for contextualization. Furthermore, a system can use all available information displays that can be perceived by the user. For example, a system could use a loudspeaker in the room of the user, a big screen display, or a PDA of the user.

- *Adaptation Goal*: Why does the system adapt to this information? Mostly, adaptive systems adapt to their user for ergonomic or pedagogical reasons. In contextualized computing, the interaction of the user with his/her physical environment becomes more important. One goal of contextualized systems could be to create more natural and authentic interactions and to make the adaptive system appear smarter in that it is appropriately embedded in real-life experiences.
- *Adaptation Strategy*: What steps are taken to adapt the system to the user, and how active or reactive are the user and the system in the adaptation process? For contextualized computing, shorter feedback loops can be predicted for the implicit tracking of the user's behaviour, as more parts of the user's behaviour can be sensed and taken into account. Additionally, because of the variety of input and output channels, new forms of interaction and continuous updates of user and context models become more important.

Based on the different dimensions of adaptive methods and therefore adaptation we would like to define our notion of adaptation, personalization, and contextualization and how these put different focus on the parameters and values used as adaptation mean, adaptation target, they follow different adaptation goals and strategies.

- *Adaptation*: At the core of adaptive systems are adaptive methods, which take an adaptation mean as a certain user characteristic or a part of the current user context and adapt different adaptation targets to this adaptation mean (Leutner 1992). In the literature a variety of such combinations of adaptation means and adaptation target can be found (Specht 1998). Considering the process of adaptation (Benyon and Murray 1993; Kobsa and Wahlster 1989; Oppermann 1994) often the distinction between adaptivity and adaptability is essential. This distinction also has an important impact on user acceptance, motivation, and new forms of personalization that are closer to customization. Personalization and contextualization can be seen as specialized forms of adaptation as they mainly consider different forms of adaptation means and adaptation targets and while the effects of personalized instruction are researched and demonstrated in a variety of works, effects of contextualized learning support are just in early stage.
- *Personalization* allows users to obtain information and services that are adapted to their needs, goals, knowledge, interests or other personal characteristics. User models deliver the adaptation mean for selecting and adapting information presentation to the individual user. On the background of web 2.0 technologies personalization has more to be seen from a customization perspective on a higher level where users adapt their learning environments by using and integrating existing services and lower level services of adaptivity as in context-aware systems that filter or pre-select information and services available based on contextual parameters.
- *Contextualization* complements personalization so that environmental states or the context of use can also be taken into account for adaptation. The combination of approaches for user modelling and context modelling can deliver much more valid data about users and their current activities. As stated above contextualiza-

tion often refers to an adaptation on low level data that can be aggregated into higher level contextual parameters. Recent works in context-aware systems (Dey and Abowd 1999; Zimmermann et al. 2005a, b) show a variety of examples of such aggregation infrastructures and how they can enhance interaction and learning support.

Considering the development of computer based instruction personalization has been an issue of discussions in different research fields (Brusilovsky 1996; Kim 2002; Kobsa 2001, 1989). Personalized and individualized computer based instruction has already been discussed since the 1960s and even earlier (Pask 1964). Nevertheless via recent developments in social software, computer based community support, and learning in distributed learner networks personalization gains a new notion. Adaptation to the individual is more and more driven by information about the social environment as the behaviour of peers and other users in similar contexts than by information based on user or student models from the individual interaction history.

Learning Networks (reference to previous chapter) are often based on individual profiles and it became obvious in the last years that the possibility of computer based networking for learning is closely related to the social aspects of learning, motivation, collaborative and peer based learning support. Social networking platforms like LinkedIn or Facebook aim to connect people based on profile information the users enter themselves. In discussions some years ago in the area of user model acquisition actually user given information was often considered a weak approach. Through the integration of social context, reputation systems, and competitive components the users visibility and presence in a social network gained a new importance. Allowing users to explore a social network via their friends and friends of friends gives them a high motivation to give personal information and even very detailed personal profiles to those services.

The scepticism about giving personal information to an unpredictable, in-transparent personalization engine has changed towards using personal profiles for self-marketing and networking. The fact that those personal profiles today can be exploited and are exploited even more effectively seems to be not really relevant to most users of those services. The effect of getting a personal access to a huge social network seems to highly out value the fear of personal tracking and misuse of information.

In that sense personalization is an essential added value and more and more personalized services allow for personalized user portals in the Web 2.0. In the past personalization was mostly interpreted as an adaptation of a machine service to a human user. Examples for adaptive methods and personalization can be found in different research areas as Intelligent Tutoring Systems (Anderson et al. 1989; Carbonell 1970; Clancey 1987), Adaptive User Interfaces (Brusilovsky et al. 1995; Carroll 1984), Adaptive Hypermedia (Brusilovsky 1996), Intelligent Multimedia or Intelligent Agents (Feiner and McKeown 1993; Johnson et al. 1998). By developments in the area of computer based social networking, social navigation, or social recommender systems adaptation and personalization gets a new notion where the

adaptation mean becomes much more the contextual information than the information about the individual and his/her history to adapt to.

A similar development can be seen within mobile learning support: by using a variety of information about the user's environment more and more environmental parameters can be used via sensors for steering the ** adaption processes for an individual user or even user groups. In the last years new research approaches like context-aware computing (Dey and Abowd 1999; Gross and Specht 2001; Oppermann and Specht 2000; Schilit et al. 1994; Zimmermann et al. 2005b) extended the idea of personalization for not only adapting to a personal profile of a user (preferences, knowledge, or cognitive characteristics) but also the context of use (Oppermann and Specht 2000) became more and more important.

The notion of context varies across many different research areas. In general, context is understood as a knowledge type that facilitates a selective and focused processing of information. In context-aware computing a variety of notions and interpretations has developed over the years. Zimmermann et al. distinguish between definitions by synonym or definitions by example which mainly name and describe certain context parameters as location, identity, time, temperature, noise, as well as beliefs, desires, and commitments and intentions (Zimmermann et al. 2007). Furthermore they introduce an operational definition of context describing following main categories of context information:

- Individuality Context, includes information about objects and users in the real world as well as information about groups and the attributes or properties the members have in common.
- Time Context, basically this dimension ranges from simple points in time to ranges, intervals and a complete history of entities.
- Locations Context, are divided into quantitative and qualitative location models, which allow to work with absolute and relative positions.
- Activity Context, reflects the entities goals, tasks, and actions.
- Relations Context, captures the relation an entity has established to other entities, and describes social, functional, and compositional relationships.

Examples for mobile systems using contextual parameters to adapt the learning support have been analysed in DeJong et al. (2008). In this work we have developed a reference model for mobile social software describing the adaptation in those systems based on a modification of the dimensions presented earlier as adaptation mean, adaptation target, adaptation goal, and adaptation strategy (see Sect. 14.3).

Several projects considered in the review have looked at how to use and enrich contents with contextual metadata (Equator Project 2003; Specht and Kravcik 2006). Mostly interesting new approaches in context-aware systems see the main strength in combining different context parameters for user support. Even more new approaches tend to combine different forms of metadata about learning objects and media to allow for flexible and contextualized learning support. In the MACE project the combination of various types of content, usage, social and contex-

tual metadata enables users to develop multiple perspectives and navigation paths that effectively lead to experience multiplication for the learner himself (Stefaner et al. 2007).

Identity context is also often combined with other forms of context. One specific example of such a combination is given by Ogata and Yano (2004b) who present CLUE, a system for learning English in real-world situations. CLUE uses (1) a learner profile to adapt the learning content to the learner's interest and (2) location information to link objects/locations to suitable English expressions, i.e. appropriate learning content. Likewise (Bo 2002), combines a user profile and user position, to facilitate personalised and location-based information delivery. AwarePhone (Bardram and Hansen 2004) uses several context-parameters at the same time. First of all, location is used to locate fellow employees within the hospital. Second, a calendar artefact is used to capture and share time context and also indicate the activity of a user at a certain moment. The activity is furthermore given by a shared status message. The combination of these three context parameters lead to what the writers call "context-mediated social awareness".

Environmental context information is used in several systems, most notably QueryLens (Konomi 2002) which focuses on information sharing using smart objects. Moreover, the TANGO system presented by Ogata and Yano (2004a) and the Musex system (Yatani et al. 2004) detect objects in the vicinity by using RFID tags. Moop (Mattila and Fordel 2005) couples a GPS location to observations/information gathered in the field for later analysis in the classroom. Wallop (Farnham et al. 2004) allows its users to discover social relationships and provides social awareness by storing and analysing social context information; to derive the social context communication patterns system interactions and co-occurrence information were analysed.

In the effort to create added value in ubiquitous environments, the current context of use becomes more important as the computer disappears. From our point of view, it is not enough to supply content or services that consider single environmental or user characteristics; we need to identify approaches for the integration and interpretation of different sensing components for modelling the user and the context more appropriately. Jameson (2001) describes how modelling the user and modelling the context of the user can both lead to more valid assumptions about the current needs and goals of a user.

14.4 Contextualization Complements Personalisation

In the following part we would like to stress the importance of a complementary use of personalization and contextualization. This could be demonstrated by new possibilities of extending adaptation mean, target, goal, and strategy.

Adaptation Mean: As main new additional sources of information sensors can measure a variety of properties of the learners and their environments. Sensors quantify temperature, humidity, pressure, sound, lightness, magnetism, acceleration,

force, and many other properties of the environment. Furthermore, sensors measure the learner's blood pressure, body temperature or eye movements as well as the location of the learner in- and outdoor. Taking into account that this information is not only available for a single user, but for all users participating in a learning network the new possibilities for adaptation become evident. Beside the possibilities of extending adaptive methods from the lowest level of biofeedback up to the highest level of user patterns on certain contextual parameters the relation between the different participants and the social dynamics enable new forms of adaptive indicators for contextual information. Aggregating the information in a sensible way for supporting contextualized learning is a key issue in generating learning network services, several examples for using footprint information, location information, object interaction history, or metadata about user processes are described in the following chapter of this section.

Adaptation Target: An adaptation of these application parts may be immediately visible or noticeable to the learner or the adaptation effects may be hidden from the learner and display their impact at a later point in time. Contextualized learning applications consider three properties or parts of the system that can be tailored to the context and the learner:

1. Human-Computer Interaction, i.e. the modality needed to enter commands or data, and receive information and services. Here new forms of interaction based on research of sense-based interaction enable new forms of interacting with educational media (Benford et al. 2005).
2. Information Presentation, i.e. the methods and coding required for receiving and displaying content (front-end). Here a variety of real life and digital indicators can be considered. Special focus lies here on the distribution of indicator environments between personal and ubiquitous indicators.
3. Functionality, i.e. the features needed to perform tasks (back-end) and the ability of the application to solve a single task or a set of tasks. A main step achieved through the extension of such adaptation targets to the physical environment can be seen in performing learning activities in real life contexts and accessing the necessary functionality to do so.

Examples for adaptation targets taken from the domain of contextualized learning application can be the pace of the instruction (Leutner 1992) (Tennyson and Christensen 1988) that can be modified based on diagnostic modules embedded in the learning process or adaptation of content presentations, the sequencing of contents and others. Taking into account contextual information enables new interesting applications as the selection of information based on locations or even the streaming of media to a mobile user in changing contexts.

Adaptation Goals: As mobile interaction has stressed the problems of HCI with limited user interfaces the role of adaptation or better contextualization gained more importance. On the other hand the mobile use of interactive devices also extends the possibilities of actively constructing and knowledge resources and information in authentic contexts and collaborate remotely in authentic learning situations.

Table 14.1 Dimensions of adaptation and extensions in contextualized learning

Adaptation target	Adaptation mean	Adaptation goal
Learning goal: content, teaching method, teaching style, media selection, sequencing, time constraints, help Presentation: hiding, dimming, annotation	Learner: preferences, usage, knowledge, interests, goals, task complexity	Didactical reasons: preference model, compensation of deficits, reduction of deficits Ergonomic reasons, efficiency, effectiveness, acceptance
Extension in contextualization		
Presentation: 3D sound overlays, multimodal interfaces, augmented reality interfaces	Context sensors: location, time, lighting, noise, other user's locations, available infrastructure	Authenticity of learning: situated collaboration, active construction of knowledge, awareness

An overview of the different adaptation dimensions and their extensions via contextualization can be seen in Table 14.1.

14.5 Conclusion

The chapter has illustrated the underlying principles of personalized and contextualized learning networks with the concept of adaptive methods and their parameters. As adaptation mean a system can focus on intra-individual, inter-individual differences or environmental changes and differences. Contextualized learning networks embed learner interaction in authentic learning contexts. Besides the embedding in a community, the social context, environmental context, relation context, and time context play an important role in first approaches and applications demonstrated in the following chapters.

In the adaptation process personalization and contextualization differ mostly in the form triggers for adaptive methods are used. In user-adaptive systems mostly users take decisions or even sometimes identify the adaptive method to be instantiated or applied. In contextualized systems the context information can support the learner on different levels from environmental indicators of peripheral information to direct guidance. In context adaptive systems the available information resources for adaptation become highly complex and that leads on the one hand to a need for more transparent user interfaces for contextualization interfaces as also the need for more semiautomatic processes of information aggregation and validation for adaptive contextualized support. In that sense it also becomes important that contextual information can not only be used for filtering and triggering adaptation but can play a role for annotation and reflection to support learning in the context of use.

In the following chapters the authors will describe different approaches and applications supporting learning networks with mobile and contextualized technology

and this introductory chapter therefore functions as a reference framework for the contextualized learning networks services presented in the following based on the background in adaptive educational methods.

References

- Anderson, J. R. et al.: Skill acquisition and the LISP tutor. *Cognitive Sci.* **13**, 467–505 (1989)
- Bardram, J. E., Hansen, T. R.: The AWARE architecture: supporting context-mediated social awareness in mobile cooperation. In *ACM Conference on Computer Supported Cooperative Work*, 2004, pp. 192–201
- Benford, S. et al.: Expected, sensed, and desired: A framework for designing sensing-based interaction. *ACM Trans. Comput.-Hum. Interact.* **12**(1), 3–30 (2005)
- Benyon, D., Murray, D.: Adaptive Systems: from intelligent tutoring to autonomous agents. *Knowledge-Based Syst.* **6**(4), 197–219 (1993)
- Bo, G.: MOBILEARN: Project final report (2002)
- Brusilovsky, P.: Methods and techniques of adaptive hypermedia. *User Modeling and User-Adapted Interaction*, Vol. 6, No. 2–3 (Kluwer, Amsterdam, 1996) pp. 87–129
- Brusilovsky, P.: Efficient techniques of adaptive hypermedia: In *Intelligent Hypertext* (Springer, Berlin, Heidelberg, New York 1997) pp. 12–27
- Brusilovsky, P. et al.: Towards adaptive learning environments: In *Herausforderungen eines globalen Informationsverbundes für die Informatik: GISI 95*(Springer-Verlag, Berlin 1995) pp. 322–329
- Carbonell, J. R.: AI in CAI: An artificial intelligence approach to computer aided instruction. *IEEE/MMS* **11**(4), 190–202 (1970)
- Carroll, J. M., Carrithers, C.: Training wheels in a user interface. *Communications of the Association for Computing Machinery* **27**, 800–806 (1984)
- Clancey, W. J.: *Knowledge Based Tutoring: The GUIDON Program* (MIT Press, Cambridge 1987) p. 377
- De Jong, T., Specht, M., Koper, R.: A reference model for mobile social software for learning. *International Journal of Continuing Engineering Education and Life-Long Learning (IJCEELL)*. **18**(1), 118–138 (2008)
- Dey, A. K., Abowd, G. D.: Towards a better understanding of context and context-awareness. *Technical Report GIT-GVU-99-22*(College of Computing, Georgia Institute of Technology 1999)
- Green, H., Hannon, C.: “*Their Space*”: *Education for a digital generation*. (2007)
- Equator Project: The EQUATOR Interdisciplinary Research Collaboration (2003)
- Farnham, S. et al.: Wallop: designing social software for co-located social networks. In *37th Annual Hawaii International Conference on System Sciences* (Hawaii 2004) pp. 107–116
- Farzan, R., Brusilovsky, P.: Social navigation support in e-learning: what are the real footprints? In *Intelligent Techniques for Web Personalisation (ITWP’05)* (Edinburgh, Scotland 2005)
- Feiner, S. K., McKeown, K. R.: Automating the generation of coordinated multimedia explanations: In *Intelligent Multimedia Interfaces* (AAAI Press/MIT Press, Boston 1993) pp. 117–138
- Gross, T., Specht, M.: Awareness in context-aware information systems. In *Proceedings of the Mensch und Computer – 1. Fachübergreifende Konferenz*, Bad Honnef, Germany (Teubner-Verlag, Bad Honnef, Germany 2001) pp. 173–182
- Heckmann, D.: *Ubiquitous User Modelling* (Akademische Verlagsgesellschaft Aka GmbH, Berlin, Germany 2006)
- Hilzensauer, W.: The MOSEP E-portfolio course: a new didactic concept for teachers and vocational trainers. *ERCIM-News* **71**(October), 18–19 (2007)
- Jameson, A.: Modelling both the context and the user. *Personal Technologies* **5**(1), 29–33 (2001)
- Johnson, L. W., Rickel, J., Stiles, R., Munro, A.: Integrating pedagogical agents in a virtual environment for training. *Presence*, **7**(6), 523–546 (1998)
- Kim, W.: Personalization: definition, status, and challenges ahead. *J. Object Technol.* **1**(1), 29–40 (2002)

- Kobsa, A., Wahlster, W. (Eds.): *User Models in Dialog Systems*. (1989), Berlin: Springer
- Kobsa, A.: Generic user modeling systems. *J. User Model. User-Adapt. Interact.* **11**(1), 49–63 (2001)
- Kobsa, A., Wahlster, W. (eds.): *User Models in Dialog Systems* (Springer, Berlin 1989)
- Konomi, S. I.: QueryLens: beyond ID-based information access. In *UbiComp 2002: Ubiquitous Computing: 4th International Conference*(Springer, Berlin, Göteborg, Sweden 2002)
- Leutner, D.: *Adaptive Lehrsysteme; Instruktionspsychologische Grundlagen und experimentelle Analysen* (Beltz, Weinheim 1992) p. 246
- Mattila, P., Fordel, T.: MOOP – using m-learning environment in primary schools. In *MLEARN 2005*, Cape Town, South Africa, 2005
- OGata, H., Yano, Y.: Knowledge awareness map for computer-supported ubiquitous language-learning. In *The 2nd IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE '04)*, 2004a, pp. 19–26
- OGata, H., Yano, Y.: Context-aware support for computer-supported ubiquitous learning. In *The 2nd IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE '04)*, 2004b, pp. 27–34
- Oppermann, R.: Adaptively supported adaptability. *Int. J. Human-Comput. Stud.* **40**, 544–472 (1994)
- Oppermann, R., Specht, M.: A context-sensitive nomadic exhibition guide. In *Proceedings of the Second Symposium on Handheld and Ubiquitous Computing* (Springer-Verlag, Bristol, UK 2000) pp. 127–142
- Pask, G.: The adaptive teaching system: In *Teaching Machines and Programming* (Pergamon Press, Oxford, London 1964)
- Schilit, B. N. et al.: Context-aware computing applications. In *Proceedings of the Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA, 1994, pp. 85–90
- Schoen, D. A.: *The Reflective Practitioner: How Professionals Think in Action*. (1983), London: Maurice Temple Smith
- Specht, M.: Adaptive Methoden in computerbasierten Lehr/Lernsystemen, *GMD Research Series* (University of Trier, Trier 1998) p. 150
- Specht, M., Kobsa, A.: Interaction of domain expertise and interface design in adaptive educational hypermedia. In *Workshop on Adaptive Hypermedia Held on Conference on UM99 and WWW99*, 1999, pp. 89–93
- Specht, M., Kravcik, M.: Authoring of learning objects in context. *Int. J. E-Learn.* **5**(1), 25–33 (2006)
- Stefaner, M. et al.: MACE – enriching architectural learning objects for experience multiplication. In *2nd European Conference on Technology Enhanced Learning EC-TEL*(Springer, New York 2007) pp. 322–336
- Tennyson, R. D., Christensen, D. L.: MAIS: an intelligent learning system, In *Instructional Designs for Microcomputer Courseware* (Erlbaum, Hillsdale, NJ 1988)
- Wild, F. et al.: First International Workshop on Mashup Personal Learning Environments (MUP-PLE08). In *3rd European Conference on Technology Enhanced Learning (EC-TEL'08)* (CEUR Workshop Proceedings, Maastricht, The Netherlands 2008)
- Yatani, K. et al.: Musex: a system for supporting children's collaborative learning in a museum with PDAs. In *The 2nd IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE '04)*, 2004, pp. 109–113
- Zimmermann, A. et al.: Personalization and context-management. User modeling and user adaptive interaction (UMUAI), Special Issue on User Modeling in Ubiquitous Computing **15**(3–4), 275–302 (2005a)
- Zimmermann, A. et al.: Applications of a context-management system. In *Proceedings of the 5th International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT-05)*(Springer, Paris, France 2005b) pp. 556–569
- Zimmermann, A. et al.: An operational definition of context. In *6th International and Interdisciplinary Conference on Modelling and Using Context (Context 07)*, Lecture Notes in Artificial Intelligence 4635, Roskilde, Denmark, August 20–24, 2007, pp. 558–571

Chapter 15

Mobile Social Software to Support Authenticity

Tim de Jong and Marcus Specht



15.1 Introduction

Recently, mobile devices have become more and more popular. Their portability and increasing possibilities to create and view high quality multimedia content make them unique tools to support learning in context. In addition, most of these devices provide ways of communication between the learner and his peers. Already a couple of years ago, Rheingold recognised the power of loosely coupled, ad hoc communities “Smartmobs”, which were based on mobile or ubiquitous access to personal social networks (Rheingold 2002). Like Rheingold, we also believe that easy access to learning content and social peers facilitates the creation of an active learning network.

T. de Jong (✉)

Centre for Learning Sciences and Technologies, Open University of the Netherlands,
Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands
e-mail: tim.dejong@ou.nl

In professional development a learner is involved in several learning activities in different contexts at the same time. In this sense, a great deal of learning is informal and therefore highly unstructured (Livingstone 2001). Mobile devices offer possibilities to make use of these spontaneous, unstructured learning situations. In addition, mobile technology should be seen as a mediating artefact (Sharples 2007) that (1) can be used to give more structure to informal learning and (2) integrate informal learning into blended learning scenarios. Koper and Tattersall (2004) support the potential of mobile devices for learning, by arguing that mobile devices offer new opportunities “to create flexible, rich and interactive learning environments”. Moreover, they specifically identify the potential of mobile information access for professional development as being able to reach anything, anyone, and anywhere. Furthermore, mobile access to personalised content provides an instant way of accessing and collecting personal memories. More specifically, mobile access to, for instance, educational blogs (Oravec 2002) can provide the learner with a way to instantly collect personal information and document learning experiences. In that way offering simple mobile tools integrated with back-end infrastructures can support long-term informal learning processes embedded in authentic contexts (Trafford 2005).

In this sense, situated learning as introduced by Wenger and Lave (1991) states the importance of knowledge acquisition in a cultural context and the integration in a community of practice. Learning in a community of practice should use authentic tasks and learning situations, i.e., settings and applications that would normally involve the knowledge learned, examples are also layed out in Chaps. 14, 16, and 17 of this section. Additionally, it should involve interaction with the social environment of the learner. Situated Learning is often contrasted with the classroom-based learning where most knowledge is out of context and presented de-contextualised. Sticht (1975), shares emphasis with situated learning in addressing the need to make learning relevant for the work context. Moreover, he states that the assessment of learning requires a context/content specific measurement. An example of the benefits of learning in a social environment is given in Sharples (2007), where children encourage each other to learn by competing with each other. The children stimulate each other to collect as much learning content as possible.

In addition to that, several challenges for collaborative infrastructures for collaborative work environments are presented in Laso-Ballesteros (2006). Among the challenges identified are: activity-oriented context-aware collaboration features provided by the collaborative infrastructure supporting human interactions, pervasive collaboration support and, heterogeneous devices with embedded collaboration capabilities. Context-aware mobile devices providing communication and collaboration features could provide a solution for these challenges (Lundin and Magnusson 2002).

In this chapter, we will first describe a number of problems encountered in current professional development approaches. To tackle these problems we will subsequently provide a systematic approach to create sound mobile scenarios from a theoretical and technological view. The approach is elucidated by examples of current and possible future solutions in mobile social software for learning.

15.2 Problems in Distributed Learning Support

Learning networks as described in Koper and Tattersall (2004) provide multi-platform access to learning content, but focus on web-based and desktop delivery of learning content. With the recent uptake of mobile devices (Castells 2007), mobile information access has become more and more important and easy to use in everyday situations. In addition, this new technology's impact on communication and learning in the younger generation is described as highly relevant for new forms of learning support (Green and Hannon 2007). However, the integration of mobile device technology and other new learning media with learning networks, such as smart phones, tablet PCs, smart boards, and gaming consoles, is mostly left out of scope. Moreover, the contextualisation of the learning content is limited. Since mobile devices offer unique possibilities for contextualised content creation and delivery, an extension with mobile devices would therefore offer the possibility to add real-world, context-specific learning support in learning networks. In that sense, the following problems with current ideas for learning network systems have become clear:

- *De-contextualisation of learning activities*: often learners have been confronted with course information without a real application context and there was often a gap in transferring knowledge to performance that could not be filled instantly by the learners. Furthermore, learning in every-day life is taking place in many occasions, only some of them formal and focused on a clear learning goal with a specified outcome, which was prepared in advance or even obvious before.
- *No support for distributed learning activities and distributed notifications*: especially, in a more activity-oriented learning paradigm the flexible and mobile support for learning activities becomes essential. Activities in this sense are combined in blended learning scenarios, which combine traditional with new technology-based learning media, and can range from reading documents, working on assessments on a PC screen, listening to a podcast, or collecting pictures on a field trip. Moreover, notifications or process reminders could be more broadly used to structure learning and draw attention to important events in the learning network or interesting aspects of the user's environment. With mobile devices these notifications can be injected and integrated in parallel multiple informal learning activity lines.
- *No Integration of Personalised and Contextualised Support for Professional Development*: informal learning, in its broadest sense, takes place everywhere, anytime and in a context or situation that is often not known beforehand. Also, it heavily depends on the learner's individual situation: her preferences, her interests, her working situation, her spare-time to study. In professional development personalised and contextualised learning should ideally be combined and tightly integrated. An integration of both personalisation and contextualisation of learning could tailor learning material to the learner's preferences and her current context at the same time (see Chap. 14).

- *No Continuous Support and Integration of Formal and Informal learning*: In the literature formal and informal learning are mostly distinguished in the sense:
 - Formal education takes place “when a teacher has the authority to determine that people designated as requiring knowledge effectively learn a curriculum taken from a pre-established body of knowledge . . . whether in the form of age-graded and bureaucratic modern school systems or elders initiating youths into traditional bodies of knowledge.” (Livingstone 2001)
 - Informal learning is “any activity involving the pursuit of understanding, knowledge or skill which occurs without the presence of externally imposed curricular criteria. Informal learning may occur in any context outside the pre-established curricula of educative institutions.” (Livingstone 2001).

Current research stresses more and more the role of supporting informal learning activities and integrating them with formal and professional development approaches in learning networks (Koper 2005). From our point of view, the role of continuous and ubiquitous support for learning activities in learning networks is essential to embed learning into every-day living, working and learning and to support situated and informal learning in learning networks.

Because of the problems encountered in on-the-spot learning, important opportunities to learn might pass. Opportunities, in which a specific situation or context could be made use of to support the learning process. To address the above problems, and with keeping social-constructivist learning theory in mind, several requirements for new mobile learning tools can be given. Mobile social software for learning should:

- enable active construction of knowledge,
- use authentic problems,
- allow for multiple perspectives in learning,
- allow for reflection about own knowledge,
- enable learning by social interaction within communities.

The requirements can be addressed by using already existing context-aware techniques, which make it possible to adapt the learning content to a certain learning moment, allow for flexible creation of media and related context information, and make the learner aware of possible situations of interest by using notifications (Dey and Abowd 1999; Oppermann and Specht 2006). Additionally, most of the current mobile learning software aims at making learning content and learning activities available to the learner anywhere and anytime. Moreover, mobile social software can be used to allow the user to communicate with learning peers, get aware of changes in the learning network and let them access multimedia information, and create documentation of learning. The next sections will provide a systematic approach to address the problems mentioned and at the same time illustrate how and to what extent the identified requirements have already been used in the current mobile learning services.

15.3 Mobile Learning Support Based on Reference Model

In an earlier paper (De Jong et al. 2008), the authors analysed the current state-of-the-art in mobile social software for learning. It became clear that a majority of the mobile learning applications today can be classified on the basis of five dimensions: content, context, information flow, pedagogical model, and purpose. These five dimensions formed the basis for a reference model for mobile social software as can be seen in Table 15.1.

As can be seen, the dimensions in the model all define different aspects of the same system, while content and information flow are relatively simple concepts, the context information used in a mobile applications, the pedagogical model, and the purpose are increasingly complex aspects. More specifically, the purpose of tool is made up of the combination of the three lower-level dimensions: content, context, and information flow. In addition, each tool has an underlying pedagogical model. Thus, the reference model of Table 15.1 can be also be perceived as a combination of:

- the information exchanged (content), the *what*,
- the information exchange between several actors (information flow, the *who*.),
- the situation this content is exchanged in (context), the *where*, and the *when*
- the pedagogical model, thus, the educational theory that forms the foundation of the mobile learning application, the way it should be designed, the *how*, and finally,
- the reason the system was made (purpose), the *why*.

The reference model was derived as a classification of existing mobile social software for learning. However, by defining the combination of different parts that make up a mobile learning solution, the reference model can be used to derive new learning solutions too. A possible way of applying the model could be the creation of innovative extensions to existing non-mobile approaches: by investigating the

Table 15.1 A reference model for mobile social software

Content	Information flow	Context	Pedagogical model	Purpose
Annotations	One-to-one	Individuality context	Behaviourist	Sharing content and knowledge
Documents	One-to-many	Time context	Cognitive	Facilitate discussion and Brainstorming
Messages	Many-to-one	Locations context	Constructivist	Social awareness
Notifications	Many-to-many	Environment or activity context Relations context	Social constructivist	Guide communication Engagement and immersion

existing approach on the basis of the reference model, new mobile solutions can be derived by analysing which parts of the reference model are already available in the existing solution, and which parts can be added in a mobile extension. All in all, the model provides us with the dimensions for creating a mobile learning solution. In the next sections, we will illustrate how parameters of these dimensions can be combined to create a mobile learning service. The examples we give will be based on already existing state-of-the-art technology, and will be described classified on their purpose:

- Sharing content and knowledge
- Facilitate discussion and brainstorming
- Social awareness
- Guide Communication
- Engagement and Immersion

Each of the following application examples will be described in a standardised way. First, an introduction is given to explain the purpose of the application in more detail. Then, the theory behind the purpose of the application is given. Then, the reasons why we think mobile support is especially suited to address the given purpose. Finally, an example of the state-of-the-art is given, which illustrates how a solution with a certain purpose can be derived on the basis of the reference model.

15.3.1 Application 1: Sharing Content and Knowledge

A main purpose found in social software systems for learning support is *sharing content and knowledge* among a community of users. Interactive Logbook presented in Bull (2004) and Chan (2005) is a shared workspace system supporting mobile learning with access to documents for mobile access and handwritten editing of content. A variety of learning management systems extend their basic functionality with services and software for mobile access (Bo 2002; Houser and Thornton 2005; Vavoula et al. 2002). Beside the sharing of content some approaches also provide facilities for the collaborative annotation of content and note taking, where personal messages can be attached to the learning content to facilitate personal or community reflection, as for example in c-notes (Milrad et al. 2002). Interestingly, the KLIV project (Brandt et al. 2002; Brandt and Hillgren 2003) delivered video content to PDAs used by nurses and demonstrated the importance of the fact that the content was created by the same user community even if the creation and the usage of content were strictly separated. QueryLens (Konomi 2002) is one example of such a system, in which a community of interest develops around real-world content, in this case music. Environmental context information is used in several systems, most notably QueryLens (Konomi 2002) which focuses on information sharing using smart objects.

Sharing content knowledge for learning is based on constructivist theory (Bruner 1966), which brings forward learning as an active process, in which learners should

construct new ideas or concepts based on their current knowledge. Learning has to take into account experiences and contexts that make the student willing and able to learn. Bruner (1996) additionally states that learning should include social and cultural aspects. Mobile devices provide unique opportunities to deliver and thus share knowledge in a wide range of situations. In addition to that, most of the devices have advanced multimedia capabilities that not only make it possible to deliver rich learning content (text, audio, pictures, video), but also ease the creation of learning content on-the-spot. We believe the ease of content access and creation also results in more active participants; and will also lead to a more active learning network with a great diversity in content.

As part of our research, we developed mobile software for contextualised language learning. The software makes it possible to learn a language by interacting with real-world objects. The learner is presented with multimedia language content coupled to real-world objects; for each real-world object or location words and sentences are given that relate to that specific learning context. Textual and audio representations are presented along with pictures explaining the meaning carried by the words and sentences. For example, in a restaurant only words and sentences will be given that relate to activities carried out in a restaurant, i.e. ordering food or drinks, while on a market place, for instance, words and phrases would be given that relate to different products that can be bought there. Hence, the software makes it possible to learn language on-the-move, in authentic situations, by using environmental or location context. The enriched objects are involved in a one-to-many information flow with the mobile devices, and the learners. The implemented scenario is a classic scenario of situated learning in a meaningful context for the learners; the association of words, sentences to the concepts (objects) they belong to. Furthermore, the scenario can be extended to allow for content creation to. In that case, the learners, can provide comments and content in their respective native languages, in that way forming a multilingual language learning network themselves. In further research we plan to investigate the effects of mobile language learning communities also creating their own content associated to real world situations. An example for such an approach will also be given in the following example.

15.3.2 Application 2: Facilitating Discussion and Brainstorming

One of the most popular classroom scenarios are mobile social tools that *facilitate discussion and brainstorming*. Mobile notes (Bollen et al. 2006) facilitates brainstorming and discussion via different kinds of annotations and voting. Ng'ambi (2005) presents a shared knowledge approach called DFAQ; a mobile social system for frequently asked questions. Yet another kind of solution aims at collaborative concept mapping, for instance PerkamII, a knowledge awareness map for sharing knowledge, collaborating and sharing individual experiences (El-Bishouty et al. 2006). PerkamII and QueryLens (Konomi 2002) additionally use real-world objects tagged with RFID to interact with real-world contexts. Mobile recommender systems like MovieLens Unplugged (Miller et al. 2003) give access to

community-generated metadata about movies and enable instant connection to real-world objects and develop a community rating and review of the described movies.

According to Cognitive Flexibility Theory (Spiro et al. 1992; Spiro and Jehng 1990), learning activities must provide multiple representations of content and support context-dependent knowledge. Especially, the theory identifies the importance of using interactive technology to support the learner in the learning process. Multiple representations can also be found in various opinions of different learners. Other research, especially in the field of knowledge management, describes the process of eliciting tacit knowledge (Nonaka and Takeuchi 1995) by contextualisation and de-contextualisation for abstraction and generalisation of knowledge. Several examples of eliciting expert's knowledge, carried out in a work context during or shortly after the actual action performed, are given by Schön (1983, 1987). Additionally, in the sense of cognitive apprenticeship (Collins et al. 1989) the learner is guided towards appropriate levels of knowledge by a constant process of contextualisation and de-contextualisation of knowledge. Cognitive apprenticeship furthermore assumes this guidance takes place in an authentic learning situation. Related to that, Wenger and Lave have stressed the importance of collaborative learning and embedding the learner in communities of practice in Wenger and Lave (1991). Additionally, Knight (2005) highlights the importance of situated learning support, by defining learning as a social practice in which learners develop their identity through participation in specific communities and practices.

Mobile social software offers a learner the opportunity to become part of a learning network and at the same time enables learning in authentic contexts. For instance, the software can be used to create *multiple perspectives on real-world objects and locations*, a user should be able to interact with a physical object and should be able to retrieve content linked to that physical object. By using shared real-world objects, multiple users can interact with them, and create information objects related to them or view, rate and comment the content added by other people (*community-generated content*). In that way, a community of users can evolve around these shared objects and the *community interaction* leads to different opinions and perspectives about these objects. The multitude of perspectives about a shared object, can lead to either a discussion between users with different opinions or leads to reflection about a situation by the learner; either by looking at the opinions of other users, or by adding content and reading it back later, as an opportunity to reflect back on what happened before (Schön 1983, 1987). To prevent the user from being overwhelmed by the amount of information available in a network contextualised search filters are used that only display the relevant information for a certain situation or context.

By combining these educational effects we developed an integrated application of mobile social software for learning, called *ContextBlogger* (De Jong et al. 2007 a, b). We applied the reference model to one specific instance of social software: weblogs. Already, some weblogs offer the creation and display of blog entries via mobile devices.¹ The ContextBlogger software combines a weblog with learning

¹www.nokia.com/lifeblog

content (*documents*) and the possibility of community-generated *annotations*, with mobile access to that content and the storage of context information. Currently, two types of context information are used and coupled to the learning content created. First, *locations context* in the form of GPS location data has been used to deliver content specific to that location. Second, the use of rectangular barcodes enables the use of *environmental context*: real-world objects are tagged with Semacodes (2006) that form a link to the learning content in the weblog.

The innovative part of the system is that it uses mobile social software and tries to form social communities by interacting with real-world objects that surround the users. To enrich the interaction with information from virtual social software communities, the real-world objects are equipped with tags to identify them. Hence, ContextBlogger provides an informal learning solution that aims at learning networks evolving around real-world objects and locations. The system addresses professional development, by providing several opportunities for the self-centred learner or a community of these learners to structure the learning process. This community creates their own context-enriched learning content and moreover should be able to view the content created by others and discuss about it by using the commenting feature of the blog. Also the system relies on the implicit assumption of lifelong learning that the responsibility for the creation and structuring of learning content resides with the self-directed learner himself (Koper and Tattersall 2004).

15.3.3 Application 3: Improving Social Awareness

Another application of mobile social software mainly aims at *social awareness*. Several different types of awareness can be identified. Nova et al. (2005) have researched the impact of location awareness on collaborative task performance. Another example for a technical solution of location-based campus support has been given in Ferscha et al. (2001). A related but more extensive system is presented in Eagle and Pentland (2005), which combines location and interest-based awareness, and should increase serendipity. In this example, location is coupled to a similarity in user profiles to notify people when someone with similar interests is nearby. Messeguer et al. (2006) describes a system that uses location information in a classroom for group awareness and identification group structure. Conversely, Kajita and Mase (2006) bring forward a classroom system that makes the teacher aware of current problems of and progress of students.

By improving social awareness a learner can more easily find his place in a learning network by making him aware of his position within this network, resulting in a more active and engaged participation in this community (Wenger and Lave 1991). Additionally, social awareness can help a learner spot interesting learning opportunities, become aware of learning peers, and plan learning opportunities to not interfere with other of activities of the learning peers.

AwarePhone (Bardram and Hansen 2004), a system for activity awareness in a hospital. AwarePhone displays the current activity of a doctor and can be used by nurses to find out an appropriate moment to interrupt him or her. AwarePhone

makes communication between peers possible by direct phone communication or sending messages. AwarePhone uses several context-parameters at the same time. First of all, location is used to locate fellow employees within the hospital. Second, a calendar artefact is used to capture and share time context and also indicate the activity of a user at a certain moment. The activity is furthermore given by a shared status message. The combination of these three context parameters lead to what the writers call “context-mediated social awareness”.

15.3.4 Application 4: Guiding Communication

Another application is *guiding communication* and provide a central place for it. An example is given by Farnham et al. (2004) with their Wallop system designed for maintaining and extending of a user’s social network; it provides a central place for people to contact each other. Additionally, most of these communication systems try to bridge the gap between geographically dispersed people to include them in collaborative processes. Raymond et al. (2005) present a synchronous communication application (instant messaging, audio and video conferencing) that also provides collaboration facilities like a white board. Similarly, the HandLeR project (Sharples et al. 2002) offers conversation between mobile learners to support knowledge sharing between them and teachers, experts and peers.

Pask (1975) stressed the importance of communication, or conversations, for learning. Conversations, he emphasised, were a vehicle for making knowledge about a certain subject more explicit. Recently, Sharples et al. (2006) in referred to Pask in their theory of mobile learning to furthermore portray the importance of conversations for learning. Conversations, they argue, are a key element for constructing knowledge in collaborative tasks. Mobile technology mediating these conversations would, they argue, also improve learning. Guiding communication between learning peers also is necessary for collaborative learning and conversations are the foundation of communities of practice (Wenger and Lave 1991).

An example is given by Farnham et al. (2004) with their Wallop system designed for maintaining and extending of a user’s social network; it provides a central place for people to contact each other. Wallop makes it possible to exchange messages with social peers. More importantly, it aims at rediscovering lost contact or meeting new ones. To introduce contacts to each other, it uses the social context of a person, his peers, and their subsequent peers. Therefore, it specifically aims at increasing many-to-many information flows between these peers, based on social constructivist theory.

15.3.5 Application 5: Increasing Engagement and Immersion

Mobile collaborative games are good examples for enhanced *engagement and immersion* of users and learners by embedding in real-world contexts. In educa-

tion, games are mostly used to motivate students and to increase their participation (Mitchell and Savill-Smith 2004). Locatory is a mobile multiuser game (Unger 2005; Zielke 2005) which implements a ubiquitous version of a memory game in which two competing teams, each consisting of two collaborating people, should find the tiles corresponding to each other and which are spread throughout an office environment. In human pacman (Cheok et al. 2004), the participants play the characters of several pacmen and ghosts, and have to collaborate to obtain virtual cookies or capture pacmen, respectively.

Paiget (1970) emphasises that learning should take place with activities or in situations that engage the learners and require adaptation. Teaching methods should be used that actively involve students and present challenges to the learner. An example of such an educational game that actively involves students is Environmental Detectives which has groups of students investigating a simulated pollution scenario by combining real-world and virtual data. Environmental Detectives (Klopfer et al. 2002) combine content creation, with the storage of location context for the content created; students take pictures in an outside setting to enhance the learning experience in remote participation. Example for effects on motivation, engagement, and immersion are also given in Chap. 16, where not only participants on a field trip but as seen from the results even the integration of a remote mobile learning group with a classroom activity significantly increased interest of learners for a topic.

Table 15.2 An overview of five different application types

Application and purpose	Example	Content	Information flow	Context	Pedagogical model
Application 1: <i>Sharing content and knowledge</i>	Contextualised language learning	Documents	One-to-many	Locations and environmental	Constructivist
Application 2: <i>Facilitating discussion and brainstorming</i>	ContextBlogger	Documents and annotations	Many-to-many	Locations and environmental	Social constructivist
Application 3: <i>Improving social awareness</i>	AwarePhone	Messages and notifications	One-to-one and many-to-many	Individuality, time, activity and locations	Social constructivist
Application 4: <i>Guiding communication</i>	Wallop	Annotations, documents and messages	Many-to-many	Individuality, locations and relations	Social constructivist
Application 5: <i>Increasing engagement and immersion</i>	Environmental detectives, RAFT	Documents	Many-to-one and many-to-many	Locations	Social constructivist

15.3.6 Comparing Applications Types

In the previous paragraphs five application examples of mobile social software for learning were given. Each of these examples combined a different set of values of the dimensions of the reference model. For a better overview and comparison of the applications, Table 15.2 provides a summary of each application and the values for each dimension they combine.

15.4 A Technical Framework for Contextualized Learning Support

Most of the application examples given before are unrelated projects that each have come up with their own software architectures. Instead of duplicating the effort of implementing similar infrastructures, a standardised extendable architecture would be beneficial to a larger integration in modern-day teaching. In this sense, the technology in particular can provide an insurmountable hurdle: the market contains lots of different devices without much standardisation, which leads to a need for detailed technical knowledge to be able to integrate mobile technology in existing learning scenarios (blended learning designs). Moreover, the rapidly changing mobile technologies form an additional burden to keep the mobile learning scenarios up-to-date; even worse, while most mobile learning designs would remain the same and would need similar functionality, this would have to be implemented again and again for new technology. Last, small-scale experiments could be used to create enthusiasm and show the benefits of mobile learning to teachers, learners, and institutions. The creation of such experiments calls for flexible and fast prototyping, and by giving the opportunity to create and integrate mobile learning fast and without too much effort, the number of applications would increase, making room for new and innovative mobile learning approaches.

To address the issues above, we propose a technical framework for mobile and contextualised learning, based on existing standards that should take into account the following requirements:

- *independence of mobile technology*, which allows for a more heterogeneous user group and to some extent circumvents the demands of rapidly changing/aging technology,
- *standardised approaches for distributed technology and interaction between devices*: a client-server architecture adhering to existing web service standards. Standardised interfaces also make it possible for different technologies (smartphones, iPods, desktops, smartboards) to communicate in a standardised and similar way,
- *adhere and make use of as much technology standards as possible*: this will make the integration with existing learning management and other systems easier,

- *easy addition of new technologies* by creating new, light-weight client technology that uses existing server functionality and combines it new ways,
- *a range of commonly needed functionality*, which does not need to change frequently: instead of writing different clients for each mobile with similar functionality, create light-weight clients and group all similar functionality on a server, which calls for,
- *a modular server architecture*, in which new functionality can easily be added, and integrated within existing learning designs.

A technical framework based on these requirements should provide a system that integrates the use of content with the use of metadata, should make it possible to combine different kinds of context information into higher level information, and also should enable the design of higher level processes based on context information and available content. Additionally, the technical framework should be founded on the reference model presented earlier. Figure 15.1 gives an overview of our technical framework based on the requirements just formulated. The technical framework is comprised of a multi-column model with four layers. On the one hand, the columns identify the different kinds of properties that can be used in a learning process: the context metadata identifying the learning situation, the electronic media used in the learning process (context and content in the reference model), and the physical

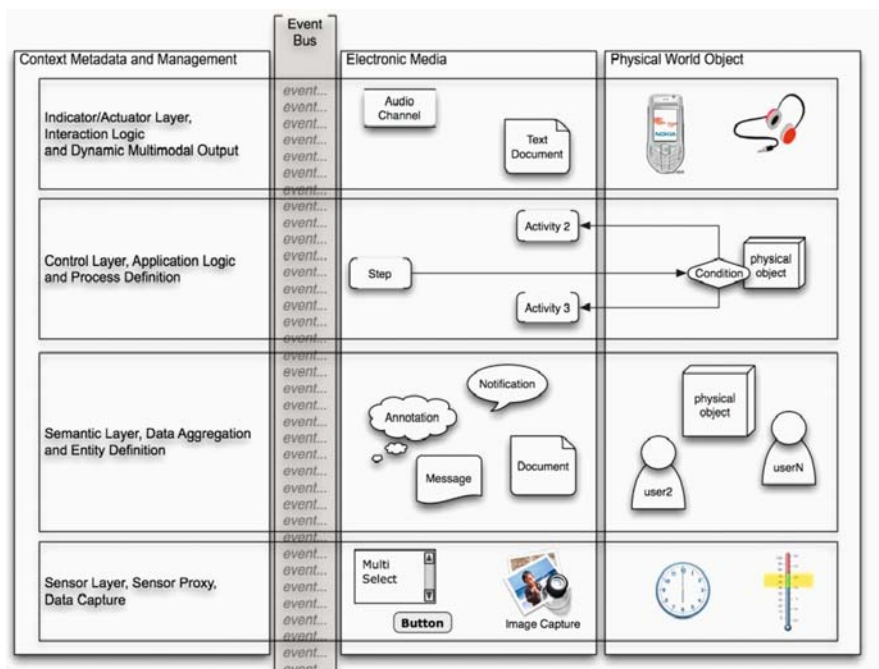


Fig. 15.1 The contextualised media framework, its layers and entities

world objects the learners interact with during that learning process. On the other hand, the four layers represent the several forms of data used in the system. For the layered model we used the infrastructure described in Zimmermann et al. (2005) as a guideline, in which in every layer the data is semantically enriched the data step by step. Based on this framework the different components as sensors, aggregators, control structures, and rendering components for different indicators and multimodal output have been implemented for the applications explained before.

The enrichment starts at the first and lowest layer, which represents the simplest form of data, electronic media information or real-world objects. The layer furthermore collects the data captured by the sensors about those real-world objects and it acquires the electronic media created by the users. The second layer groups the sensor data and electronic media into higher-level concepts that can be used to represent the real-world objects or information attached to these objects. After that, the third layer provides us with the means to define activities, define application logic and processes, and combine the context metadata to take higher order decisions on the basis of semantically enriched data (from layer two). In this layer the educational processes based on the pedagogical paradigm in the reference model can be defined. Furthermore, the information flows and conditions for the delivery of content or notifications are defined here. Moreover, the adaptation to the user's personal preferences or physical objects the user interacts with, happen in the third layer. Finally, the fourth layer carries out actions and delivers the electronic media based upon the decisions that have been taken in layer three. This layer also chooses the correct actuator and suitable content for a certain situation. For example, if the noise level is too high for people to hear an audio feedback the layer could decide to provide visual feedback instead. In short, the purpose of this layer is to carry out an action or change a real-world situation that is given by the last column.

A number of tools should expose the functionality of the technical framework to a variety of users, including technology developers but also non-technical end users. The tools should make it as easy as possible to create mobile learning scenarios using the standard functionality of the technical framework. The tools need to make it possible to create such a scenario on the basis of an application description much like the examples give before. To this cause, these tools should facilitate the combination the various components and functionalities of the technical framework in a flexible way. On the one hand, for a developer it should be possible to design new context aggregations from several easier context dimensions. The aggregation of time and changes in location, for instance, could be used to calculate the speed of a person. The designed constructs should be reusable for other users and more importantly, by combining several low-level concepts into higher-level activities or situations, concepts can be created that could also be used by a non-technical person: "if a person is speeding, please notify him that he should break".

On the other hand, a person developing the pedagogical scenario should not be bothered with these technical details; instead, this person would work on an activity- or situation-based perspective that portrays the activities/situations that make up the mobile learning scenario. The combination of context and content would take place at this level. Mobile learning scenarios that prove useful could be saved as templates

for future use. Additionally, specific tools for specific target groups should allow for easy creation of common mobile learning scenarios. For instance, a program for designing mobile language learning could be used to create scenarios ranging from daily vocabulary list shown on a mobile, to useful phrases in a foreign language adapted to a specific location (Ogata and Yano 2004).

15.5 Conclusion

In this chapter, we described a systematic analysis of already existing mobile learning solutions. The analysis was based on reference model for mobile social software for learning which was presented earlier in De Jong et al. (2008). The reference model provided five dimensions that, when combined, form the basis of a mobile learning scenario; different values of the following dimensions can be combined: content, context, information flow, pedagogical model, and purpose. The applications we described on the basis of the purpose dimensions. In that way, we illustrated how mobile learning scenarios could be created on the basis of examples encountered most in the current state-of-the-art. In an analogous way, extensions to already existing learning services or new mobile learning services could be created. Finally, we described a technical framework that was founded on the reference model, and that can be used to implement the mobile learning scenarios that were derived. By providing a standardised technical framework for contextualised learning, we provide independence of mobile technology, standardised approaches for distributed technology and interaction between devices; make use of as much technology standards for easier integration of existing software, extensible and modular server architecture. Such a technical framework furthermore facilitates easy prototyping and a minimised effort to implement new scenarios. Hence, with a standardised technical framework, like the one presented here, we hope to reduce the threshold for the implementation of mobile learning services, which will with a bit of luck lead to a more widespread acceptance and application of mobile devices in current learning support. More specifically, we believe a more extensive use of mobile devices will lead to a higher authenticity in learning, which subsequently could result in more applied and improved learning.

References

- Bardram, J. E., Hansen, T. R.: The AWARE architecture: supporting context-mediated social awareness in mobile cooperation. In *ACM Conference on Computer Supported Cooperative Work* (2004) pp. 192–201
- Bo, G.: MOBILEARN: Project final report. (2002)
- Bollen, L. et al.: PDAs as input devices in brainstorming and creative discussions. In *Fourth IEEE International Workshop on Wireless, Mobile and Ubiquitous Technology in Education (WMTE')* (2006) pp. 137–141
- Brandt, E., Björgvinsson, E., Hillgren, P., Bergqvist, V., Emilson, M.: *PDA's, Barcodes and Video-films for Continuous Learning at an Intensive Care Unit*, **10**(2), (pp. 293–294), (2002), Aarhus, Denmark

- Brandt, E., Hillgren, P., Attewell, J., Savill-Smith, C.: *Self-produced video to augment peer-to-peer learning* (pp. 27–34), (2003), London, UK: Learning and Skills Development Agency
- Bruner, J.: *Toward a Theory of Instruction* (Harvard University Press, Cambridge, MA 1966)
- Bruner, J.: *The Culture of Education* (Harvard University Press, Cambridge, MA 1996)
- Bull, S., Bridgefoot, L., Corlett, D., Kiddie, P., Marianczak, T., Mistry, C., et al.: *Interactive Logbook: the development of an application to enhance and facilitate collaborative working within groups in higher education* (pp. 39–42), (2004), Rome, Italy: Learning and Skills Development Agency
- Castells, M. et al.: *Mobile Communication and Society: A Global Perspective* (MIT Press, Cambridge, MA 2007) p. 331
- Chan, T., Corlett, D., Sharples, M., Ting, J., Westmancott, O.: Developing Interactive Logbook: A Personal Learning Environment. In *IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE'05)* (pp. 73–75), (2005), IEEE Computer Society Press
- Cheok, A. D. et al.: Human Pacman: a mobile, wide-area entertainment system based on physical, social, and ubiquitous computing. *Personal Ubiquitous Comput.* **8**(2), 71–81 (2004)
- Collins, A. et al.: Cognitive apprenticeship: teaching the craft of reading, writing, and mathematics: In *Knowing, Learning and Instruction* (Lawrence Erlbaum Associates, Hillsdale, NJ 1989) pp. 453–494
- De Jong, T. et al.: Campus memories: learning with contextualised blogging. In *The 2nd TenCompetence Workshop*, Manchester, United Kingdom (2007a) pp. 59–67
- De Jong, T. et al.: ContextBlogger: learning by blogging in the real world, In *ePortfolio 2007*, Maastricht, The Netherlands (2007b)
- De Jong, T. et al.: A reference model for mobile social software for learning. *IJCEELL.* **18**(1), 118–138 (2008)
- Dey, A. K., Abowd, G. D.: *Towards a Better Understanding of Context and Context-Awareness* (College of Computing, Georgia Institute of Technology, Atlanta, GA 1999)
- Eagle, N., Pentland, A.: Social serendipity: mobilizing social software. *Pervasive Comput. IEEE* **4**(2), 28–34 (2005)
- El-Bishouty, M. M. et al.: Learner-space knowledge awareness map in computer supported ubiquitous learning. In *Fourth IEEE International Workshop on Wireless, Mobile and Ubiquitous Technology in Education (WMTE'06)* (2006) pp. 116–120
- Farnham, S. et al.: Wallop: designing social software for co-located social networks. In *37th Annual Hawaii International Conference on System Sciences*, Hawaii (2004) pp. 107–116
- Ferscha, A. et al.: Location awareness in community wireless LANs. In *Informatik 2001*, Vienna, Austria (2001) pp. 190–195
- Green, H., Hannon, C.: *Young People Are Spending Their Time in a Space Which Adults Find Difficult to Supervise or Understand* (Demos, London 2007)
- Houser, C., Thornton, P.: Poodle: a course-management system for mobile phones. In *IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE'05)* (pp. 211–215), (2005)
- Kajita, S., Mase, K.: uClassroom: expanding awareness in classroom to ubiquitous teaching and learning. In *Fourth IEEE International Workshop on Wireless, Mobile and Ubiquitous Technology in Education (WMTE'06)* (2006) pp. 161–163
- Klopper, E. et al.: Environmental detectives: PDAs as a window into a virtual simulated world. In *IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE'02)* (2002) pp. 95–98
- Knight, S.: Innovative practise with e-Learning (JISC 2005)
- Konomi, S. I.: QueryLens: beyond ID-based information access. In *UbiComp 2002: Ubiquitous Computing: 4th International Conference*, Göteborg, Sweden (Springer, Berlin, Göteborg, Sweden 2002) pp. 210–218
- Koper, R.: TENCompetence: building the European Network for lifelong competence development (Open Universiteit Nederland, Heerlen 2005)

- Koper, R., Tattersall, C.: New directions for lifelong learning using network technologies. *Br. J. Educ. Technol.* **35**(6), 689–700 (2004)
- Laso-Ballesteros, I.: Research perspectives on collaborative infrastructures for collaborative work environments. In *15th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE'06)* (2006) pp. 3–10
- Livingstone, D. W.: Adults' informal learning: definitions, findings, gaps and future research. Working paper (2001)
- Lundin, J., Magnusson, M.: Walking and talking-sharing best practice. *IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE'02)* (2002) pp 71–79
- Messeguer, R. et al.: Group-awareness for mobile cooperative learning. In *Fourth IEEE International Workshop on Wireless, Mobile and Ubiquitous Technology in Education – (WMTE'06)* (2006) pp. 74–81
- Miller, B. N. et al.: MovieLens unplugged: experiences with a recommender system on four mobile devices. In *Human-Computer Interaction Conference HCI 2003* (Bath, England 2003)
- Milrad, M. et al.: C-notes: designing a mobile and wireless application to support collaborative knowledge building. In *IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE')* (2002) pp. 117–120
- Mitchell, A., Savill-Smith, C.: *The Use of Computer and Video Games for Learning* (Learning and Skills Development Agency, London 2004)
- Ng'ambi, D.: Mobile dynamic frequently asked questions (DFAQ) for student and learning support. In *MLEARN 2005*, Cape Town, South Africa (2005) pp. 116–119
- Nonaka, I., Takeuchi, H.: *The Knowledge Creating Company* (Oxford University Press, USA 1995)
- Nova, N. et al.: 'Location is not enough!': an empirical study of location-awareness in mobile collaboration. In *IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE'05)*, Tokushima, Japan (2005) pp. 21–28
- Ogata, H., Yano, Y.: Context-aware support for computer-supported ubiquitous learning. *The 2nd IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE'04)* (2004) pp. 27–34
- Oppermann, R., Specht, M.: Situated learning in the process of work: In *Engaged Learning with Emerging Technologies* (Springer Verlag, Berlin 2006) pp. 69–89
- Oravec, J.: Bookmarking the World: weblog applications in education. *J. Adolescent Adult Literacy* **45**(7), 616–621 (2002)
- Paiget, J.: *The Science of Education and the Psychology of the Child* (Grossman, New York 1970)
- Pask, G.: *Conversation, Cognition, and Learning* (Elsevier, New York 1975)
- Raymond, D. et al.: Bringing mobility to synchronous collaborative activities: recent enhancements of the "platine" platform. In *IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE'05)* (2005) pp. 59–61
- Rheingold, H.: *Smart Mobs: The Next Social Revolution* (Perseus Books Group, New York 2002)
- Schön, D. A.: *The Reflective Practitioner: How Professionals Think in Action* (Maurice Temple Smith, London 1983)
- Schön, D. A.: *Educating the Reflective Practitioner* (Jossey-Bass, San Francisco 1987)
- Semacode: Semacode | URL barcodes | practical ubiquitous computing. (9 November 2006)
- Sharples, M.: *Big Issues in Mobile Learning: Report of a Workshop by the Kaleidoscope Network of Excellence Mobile Learning Initiative* (Learning Sciences Research Institute, University of Nottingham 2007)
- Sharples, M. et al.: The design and implementation of a mobile learning resource. *Personal Ubiquitous Comput.* **6**(3), 220–234 (2002)
- Sharples, M. et al.: A theory of learning for the mobile age (forthcoming): In *The Handbook of E-learning Research*, London (2006)
- Spiro, R. J. et al.: Cognitive flexibility, constructivism and hypertext: random access instruction for advanced knowledge acquisition in ill-structured domains.: In *Constructivism and the Technology of Instruction* (Erlbaum, Hillsdale, NJ 1992)

- Spiro, R. J., Jehng, J.: Cognitive flexibility and hypertext: theory and technology for the non-linear and multidimensional traversal of complex subject matter In *Cognition, Education, and Multimedia* (Erlbaum, Hillsdale, NJ 1990)
- Sticht, T. G.: Applications of the audread model to reading evaluation and instruction: In *Theory and Practice of Early Reading* (Erlbaum, Hillsdale, NJ 1975)
- Trafford, P.: *Mobile Blogs, Personal Reflections and Learning Environments. ARIADNE, Web Magazine for Information Professionals in Archives, Libraries and Museums*, Vol. 44 (2005)
- Unger, M.: *Entwicklung eines kontextbasierten mobilen Spiels auf der Basis von Macromedia Flash* (FH Offenburg, Fraunhofer Institut für Angewandte Informationstechnik (FIT) 2005)
- Vavoula, G., Sharples, M., Milrad, M., Hoppe, U., Kinshuk.: *KLeOS: A personal, mobile Knowledge and learning organisation system* (pp. 152–155), (2002), Vaxjo, Sweden: The Printing House of the IEEE Computer Society & IEEE
- Wenger, E., Lave, J.: *Situated Learning: Legitimate Peripheral Participation* (Cambridge University Press, Cambridge, NY 1991)
- Zielke, L.: *Entwicklung und Evaluation eines kontextbasierten mobilen Spiels unter Einbeziehung verschiedener Lokalisierungsverfahren* (Fachhochschule Oldenburg, Wilhelmshaven 2005)
- Zimmermann, A. et al.: Personalization and context management. *User Model. User-Adapted Interact.* **15**, 275–302 (2005)

and time constraints limit the experiential elements of authentic learning settings. For example, often this experience is either overlaid with activities that are carried out in isolation to the learning setting, or it is so distracting for students it is difficult to obtain the learning returns that the experiences have the potential to generate. Information and Communication Technologies (ICT) have the potential to make the links between learning environments external from classroom settings, and facilitate authentic and situated approaches to learning not readily available except by direct and extended visitation through learning networks. Learning networks can be thought of as communities of learners who are helping each other to better understand concepts that are the focus of the particular network. Such networks can be formal or informal, but in this chapter section we are looking at formal learning networks.

This concept is not new and there have been many examples of progressive educators developing such learning settings by making use of technologies. Ambron and Hooper (1988) outlined a range of early ideas on the use of interactive multimedia that were to revolutionise thinking about the use of technology in education and professional development. Projects such as *The Voyage of the MIMI*, developed by the Bank Street College (Gibbson and Hooper 1993) and the early work on the *Apple Classroom of Tomorrow* on the *Wireless Coyote* project (Grant 1993) demonstrated effective implementations of these ideas through authentic and situated approaches to learning.

More recently web and service technologies have increased the affordances available for designing learning settings as learning networks and more sophisticated pedagogies have been able to be implemented. Projects such as the *Remotely Accessed Field Trips (RAFT)* project (Terrenghi et al. 2004), the *Sydney Olympic Park* project (Brickell and Herrington 2004; Brickell et al. 2005a, b) and the *Killalea* project are all fine examples of linking learners and professionals to authentic experiences and contexts set within locations beyond school and training settings. This paper reviews these learning settings and the learning outcomes achieved to this point, and also outlines the lessons learnt from the design and implementation of these types of authentic learning networks.

16.2 Authentic and Situated Approaches to Learning

How real does a learning environment need to be to ensure quality learning outcomes? This is a question that is often asked in relation to the use of virtual learning resources, especially when they might be used as a substitute for field trips and visitations. In setting realistic and authentic tasks, some authors argue that only a real problem situation should be used. For example, Savery and Duffy (1996) argued that in creating complex and authentic contexts for learning, the problems to be investigated must raise the concepts and principles relevant to the content domain, and importantly, the problems must be real. They stated:

There are three reasons why the problems must address real issues. First, because the students are open to explore all dimensions of the problem there is real difficulty of creating a rich problem with a consistent set of information. Second, real problems tend to engage learners more—there is a larger context of familiarity with the problem. Finally, students want to know the outcome of the problem—what is being done about the flood, did AT&T buy NCR, what was the problem with the patient? These outcomes are not possible with artificial problems. (Savery and Duffy 1996, p. 144)

Often, field experiences and visitations do enable real problem situations to be engaged by students. In such situations, data is collected and used to create genuine and viable solutions to real problems. However, not all field experiences address real problems but instead employ data collection and analysis exercises that facilitate the learning of key principles and content, rather than complex problem solving linked to solutions. Because a learning experience is centered beyond the classroom in a realistic setting such as the natural environment or museum, it does not mean it is authentic. Conversely, a learning experience can be authentic even if it does not take place in a real situation. Real experiences with maximum fidelity are not essential to truly engage students in authentic learning (Herrington et al. 2007; Alessi 1988). The physical reality of the learning situation is of less importance than the characteristics of the task design—the “cognitive realism” (Smith 1986) of the task—and the level of engagement of students in the learning environment.

In this paper, we argue that the completion of sustained and complex tasks provide a key rationale for authentic learning, rather than physical reality or verisimilitude. Ten characteristics of authentic tasks and activities have been developed from a review of research on authentic learning environments and from the literature on situated learning, anchored instruction and problem-based learning (cf. Herrington et al. 2004). Authentic activities: have real-world relevance; are ill-defined, requiring students to define the tasks and sub-tasks needed to complete the activity; comprise complex tasks to be investigated by students over a sustained period of time; provide the opportunity for students to examine the task from different perspectives, using a variety of resources; provide the opportunity to collaborate; provide the opportunity to reflect; can be integrated and applied across different subject areas and lead beyond domain-specific outcomes; are seamlessly integrated with assessment; create polished products valuable in their own right rather than as preparation for something else; allow competing solutions and diversity of outcomes (cf. Herrington et al. 2007, 2003).

Tasks that encompass these ten characteristics create experiences for students that are “cognitively real”, and the following sections describe several learning networks that exemplify this authentic approach.

16.2.1 The Remotely Accessed Field Trips Project (RAFT)

The Remote Access to Field Trips project, (RAFT) (Terrenghi et al. 2004) was a three year project supported by the European Commission that developed technologies and educational models for supporting field trip and field-to-class communi-

cation with mobile technology. Field trips are an ideal example of an established pedagogical method that can be enhanced with computer-based tools to offer new ways of collaboration and to actively construct individual knowledge. The learners in the field can collect information and contextualize it with their own experiences and at the same time work on tasks with their peers and detect new perspectives and solutions to given problems. RAFT is an example of use of these new types of learning settings where learners are embedded or connected into real world scenarios using ICTs. In this project the core resource was a service portfolio, where students with mobile devices could access, collect, and share data during field trips in cooperation with a classroom or remote group. Small groups of students went on field trips while participating classrooms participated via the web-applications displayed on a shared classroom interface or individual widget-based applications for specific tasks. Data from the field groups (text, graphics, pictures, video) as well as video conferencing facilities, allowed for direct interaction between the field and the classroom. The cooperation and collaboration was mainly structured by tasks that could be dynamically managed as questions and problems arose in the field or in the classroom.

The main objectives of the RAFT project were:

- to demonstrate the educational benefits and technical feasibility of technology enhanced remote field trips.
- to establish extensions on current learning material standards and exchange formats for contextualization of learning material. This is combined with the embedding of learning and teaching activities in an authentic real world context and the possibility to reuse and related existing and real time collected data.
- to establish new forms of contextualized learners' collaboration with real time video conferencing and audio communication in authentic contexts.

As a first step in the project, the different phases and functional requirements for supporting live collaboration and information access during field trips were established. Field trips were held in Scotland, Slovakia, Canada and Germany in order to identify different activities in the field and in the classroom and to draw first evaluations of critical factors. Furthermore, the current use of technology was analyzed and the project introduced new out-of-the-box technology to the field trip participants. Creative design workshops were held on using technology on field trips. Through these trials, different phases for preparing the field trip, experiencing the field trip in the classroom and in the field, and the evaluation after the field trip were identified. In those phases, a variety of stakeholders and participants contributed to the field trip and took an active role in it. Out of those requirements and use cases, a model for different roles, stakeholders and communications channels was developed (Fig. 16.1).

As one main result of the project, the concept of pairs working in different roles of the learning scenarios became evident, i.e., as today's technology is often not really designed with focus on mobile usability, learners paired up in teams to com-




Fig. 16.1 An overview of roles in the field and in the classroom and the communication channels

bine device handling, reflection in action, and early discussion in the onsite working pairs (see for example scout role in Fig. 16.1).

The RAFT applications enabled different participants in a synchronous collaborative learning situation to solve common tasks and learn with different activities about a topic. From the prototyping usage of the RAFT applications, the following key activities could be considered as new qualities of contextualized learning approaches:

- *Distributed cooperative task work:* The distributed work on a task focuses the interaction and communication between the learners, and technology moves to the background when the curiosity about the given task and its exploration in physical and knowledge space become the main interest. The context in this sense is an enabling means that allows the learners to become deeply immersed in the learning subject at hand.
- *Active construction of knowledge and learning materials:* Users are much more motivated when “self made” learning material is integrated in the curriculum and they have the possibility to extend existing pre-given structures for learning. Figure 16.2 shows an example of images taken by learners integrated with a biology field trip, a communication channel and a predefined task structure.
- *Full support for field trips as a blended learning process:* The different phases of field trips are essential for successfully doing a field trip. Trainers need to specify preparation materials, distribute user roles, and define field and classroom tasks.

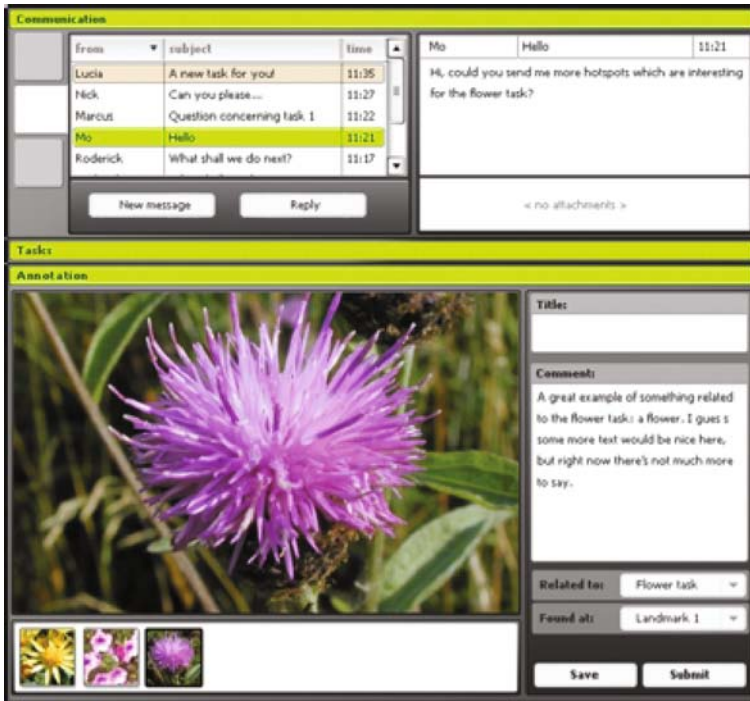


Fig. 16.2 A sample screen showing a collected image and student information

Additionally after the field trip the collected materials need to be reviewed and archived in standardized formats to ensure reuse and quality assurance. Different applications were used in different phases of the field trip support. Before the field trip started a teacher could use the RAFT Planner to define the resources and learning activities before and after the field trip, to track student progress on the learning activities, and to add additional materials and feedback on the fly. In order to define the field trip tasks, roles and communication setups, the teacher could use the RAFT Field Trip Preparation Tool. During the field trip the run time system supported cooperation and communication between the various user roles as shown in Fig. 16.1. After the field trip the authoring and reviewing took place with an Authoring Tool, which allowed users to review collected assets and learning objects and create SCORM compliant packages for later reuse.

The RAFT appliances supported the user with different tools depending on the current phase in the field trip process: preparation, field trip activity, or evaluation. Different interfaces and widgets had to be designed and implemented to give the user flexible and distributed access to the service backend and the live communication channels. During the field trip the selection of information and collaboration tools was based on the device used and current user role. Based on the experiences

of the prototyping phase the implementation of different user roles and interfaces was developed with specialized appliances. The RAFT applications can be seen as different components in a blended learning process, which were distributed in time, location, social context and across the different phases of the field trip. Furthermore the basic architecture in RAFT supported the integration of a variety of widgets using different modalities for input and output. All messages went into a common backend via a web services interface and could be used with different rendering and display widgets. This ensured the most flexibility for communicating between different interfaces in the classroom and the field.

The RAFT project evaluated field trips with more than 400 students, in about 17 schools, in a variety of curricular areas including: Art, History, Social Studies, Geography, Environmental Science, Biology and Mathematics. In two evaluation studies reported in Bergin et al. (2004) effects of video-mediated communication and of data transfer and collection on site were tested. The first study used video-mediated communication (VMC) to allow students to communicate with experts in the field. Information gained from the varying design of the two studies gave us two further insights into the potential benefits of the RAFT concept. Firstly, roles may ensure an equally high level of motivation between the classroom and field. Secondly, video-conferencing can provide a similar level of participation and interaction between field and classroom students. A second study tested students' experiences of using off-the-shelf technologies to transfer data while assigned with specific roles. Participants were school students aged 11–18 years. The findings showed RAFT resulted in high levels of interest and was both positive and engaging for students. A significant increase in student interest in topic was reported in the second study.

16.2.2 The Sydney Olympic Park Project

This Sydney Olympic Park (SOP) project was developed to offer school children access to a varied environmental site situated in Sydney Olympic Park in Sydney, Australia. The project was envisaged to offer authentic learning experiences making extensive use of technology to support student investigative processes. The long-term intention is to offer a wide coverage of curricula and topical social issues as underlying structuring of the investigations and to expose the rich resources and history of the park to learners. Set within the boundaries of SOP, the parklands of over 400 hectares of natural and remediated environments play host to a range of flora and fauna.

Work on the development of this project (Brickell and Herrington 2004; Brickell et al. 2005a, b) has utilized a situated framework in developing a model of learner engagement that views learning as an active and interpretive process of construction rather than the memorization of factual information. With its origins in a “traditional” excursion model where students are taken from the classroom environment and placed in the visited environment, albeit often with little connectivity between

the two, the approach adopted utilizes the power of ICTs to engage students through an online environment. This places the excursion in the context of an integrated three-phase process to the inquiry-based task—a pre-visit phase (classroom environment), a fieldwork phase (excursion environment) and a post-visit phase (classroom environment). The provision of a range of technology tools allows students to store, retrieve and analyze any data they collected in the fieldwork phase, and return to it at any time for further analysis. These data can then be compared and contrasted to historical data of the same measures, a facility that immediately increases the learning potential of the activities.

16.2.2.1 The Geography Challenge

The online environment was designed to engage students through the development of quality learning experiences that have relevance to practices in everyday life. Such experiences require a broad range of cognitive, linguistic and social skills that extend the intellectual capabilities of the learner leading to improved student outcomes. The specific learning sequence, represented by the learner challenge, The Geography Challenge, was developed to illustrate the challenge based approach to learning set in a specific curriculum context and aligned to the New South Wales high school geography curriculum.

On entering the online environment an animated sequence (Fig. 16.3) introduces the students to three problems faced by the Parklands management—mosquito infestation (pests), contaminated water (water management) and litter (human interaction). The “Challenge” is then presented in the form of a letter to geography consultants (the students) who are invited to investigate and advise the Parklands management on strategies to restore an ecological balance to the area.

In order to complete their challenge students, after researching the problem through the online environment, explore and carry out a series of tests through fieldwork in the Parklands using a range of tools and recording devices. Data collected is then entered into the online database for analysis before students propose a solution and write their report and recommendations to the park.



Fig. 16.3 Scenes from the introduction to the geography challenge focus on the three themes of pests, water management and human interaction

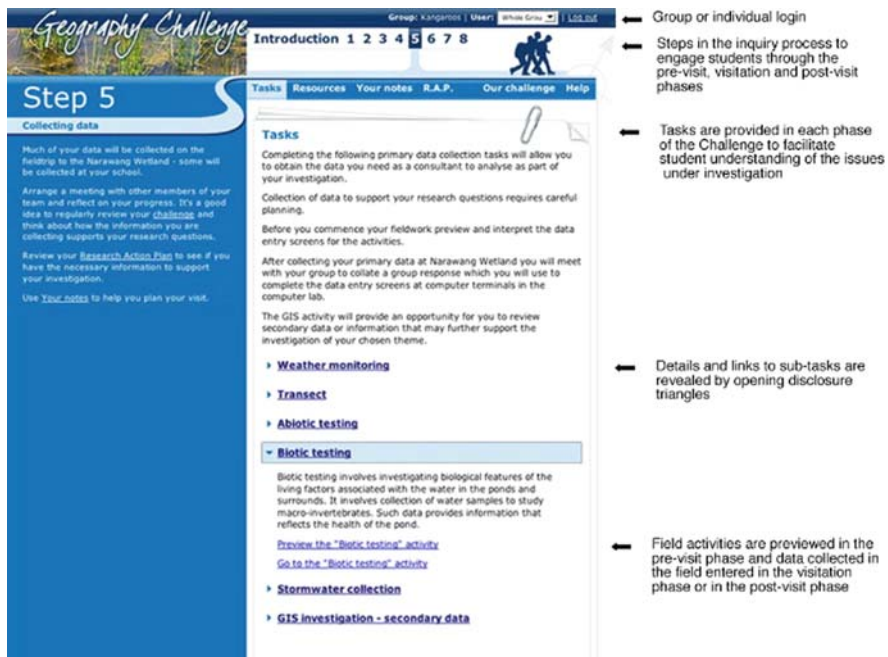


Fig. 16.4 Geography challenge interface for data collection during fieldwork

The challenge was developed with optional and essential sub-tasks in each of the three phases of student engagement. Teachers are able to guide students in choosing a set of essential tasks for pre- and post-visit experiences that are clearly integrated to the fieldwork associated with the visitation to the Parklands. Figure 16.4 shows the interface for the fifth step (out of 8 steps) in the inquiry process, the tasks for this phase of the challenge, links to the specific sub-tasks, field activities that can be previewed in the pre-visit phase at school, and links to the collected data entered in the visitation phase or in the post-visit phase of the challenge.

Usability testing was carried out with two groups of Year 10 students from local high schools—School 1 with 8 students and School 2 with 17 students—during the development. In each case students, working in pairs, were briefed about the purpose of the Geography Challenge and given general instructions on completing the tasks to allow them to explore the challenge at their own pace and in their own way. They were asked to note down any ideas or suggestions using a survey form provided (Brickell and Herrington 2007). Following the online exploration students were interviewed on both the technical aspects of the design and the pedagogical focus of the learning activities. Review of the audio transcripts have enabled the development team to address those technical concerns expressed by students. From a pedagogical viewpoint, there was general agreement that the act of placing learners in the role of professionals who are actually confronted by everyday problems, helps to extend students’ thinking to develop a broader knowledge base and strategies

to resolve the problem. Students worked together by talking about ideas and their comments included:

“You’ve got to talk it out. Each (person) had different ideas and you had to think about the consequences of that idea.”, “We swapped notes a bit and had a look at what each other said.”, “It (sharing) broadens your ideas.”

Students also commented on the effectiveness of the website in developing their knowledge of the issues associated with the Geography Challenge:

“... have a more clear understanding of an actual site (Narrawang Wetland) and increased awareness of real problems.”, “... understand geography more. It will help me gain a better understanding of wetlands.”, “... have gathered a better understanding of the wetlands and problems associated with them.”

This usability testing then formed the basis of a redesign of the learning setting with a greater emphasis on the representation of the learner activities and the role of learners during each phase of their challenge. Students believed that the web environment could support a collaborative process of problem solving and the authentic nature of the physical setting being portrayed by the web environment.

16.2.3 Killalea State Park Expeditions

The Killalea State Park is located on the East Coast of Australia, South of Sydney in New South Wales (NSW). The park has a field study centre attached which is staffed by the NSW Department of Education. The centre is strongly supported by a variety of partners and seeks to not only offer learner experiences on site, but also to offer extended learning settings for students beyond their visitation as well as those students who cannot physically access the site. The centre has had a long history of innovative access to the unique local environment, and through an innovative web environment has now been able to offer extended access to users.

The key mode of access to the site is via a visitation and then web-based resources and data are accessed to resolve the many issues that the visitation has uncovered. The Expedition project was funded through an Australian Government grant partnered with local schools and authorities. The Centre has recently been successful with a further grant to develop the learning setting.

Three expeditions, each with a residential experience, are available for class visits. The specific learner goals and the expedition problems themselves are determined by the learners. Students collect data and images through their task, enter this information at the site and then analyze this data, together with data collected by other students, to resolve their expedition goals. During the site visits, Killalea staff and young scientists recruited to the project support students.

The expedition program was designed for Years 5–8 and focused on three expeditions, Birds for Diversity, Battling Bitou Bush, and Saving Zieria. Participating students were provided with detailed briefing notes prior to their expedition, and during their site expedition experience they used a log-book to compile annotated photographs, goals they had set, surveys, investigation reports, concept maps and



Fig. 16.5 Representation of expedition challenges on the project web site

a glossary. The web site (Fig. 16.5) supported the investigation through additional data and information.

The report of the first year of operation of the project (De Landre 2007, Environmental science for schools: a middle years initiative, unpublished report) described a study of the initial implementation of the expedition project. The study surveyed 117 project participants to try to determine the effectiveness of the learner experience and the parent perception of each student's experience at the site and via the web-based environment. The student survey consisted of nine statements about the students' perceived value of the expedition experience using a five point Likert scale. An analysis of the results showed that there was an overwhelmingly positive response to the experience. For every question, over 80% of responses fell in the Strongly Agree or Agree points of the scale for statements such as:

"I was a successful member of the expedition." (91%), "I learnt things that were important to me." (87%), "My interest in science has increased", (81%) (De Landre 2007, Environmental science for schools: a middle years initiative, unpublished report, pp. 4-5).

Additionally, the parents of the participating students were surveyed to develop an understanding of parental knowledge and satisfaction with the project. A total of 66 parents of participants responded to the survey and their responses were also overwhelmingly positive. They reported that they "strongly agreed" their child's

experiences at the site and participating in the expeditions increased their child's appreciation of science (78%) and resulted in a very positive view when recounting the experience (70%). They also reported that they valued their child experiencing curriculum enrichment such as this (93%) and they recommended the continuation of and access to this program (98%).

Overall the survey response indicated extremely strong support from parents and an extremely positive view by students about the value of the learning experience. Undoubtedly, the inspiring nature of the site and the quality of the instructional support at the site did have an influence on the learner responses, but it is unusual to find such an overwhelming positive response to a learning experience for students of this age bracket.

16.3 Lessons Learnt from Learning Network Design and Implementation Process

Each of the learning networks described here attempt to locate their pedagogical approach in situated learning, anchored instruction and problem based learning. They all comply, in most respects, with the ten characteristics of authentic tasks proposed by Herrington et al. (2007, 2003) and the research outcomes indicated that learners are successful when using learning networks designed using these principles. The learning network designers learnt a number of important lessons through the design and implementation process and although the lessons were not all the same, some clear themes for design were evident.

Arguably, the most important lesson that became evident in all the learning environments described here, and one that is deeply embedded in the principles of authentic learning, is the fundamental importance of the task. The design of the task must also be carefully planned to ensure that the range of learning outcomes that must flow from its completion are fully explored and enabled. If students are given complex and realistic tasks, the chances of successful and deep learning are greatly increased. Such tasks cannot be described as authentic if they can be completed in a minimum amount of time, such as within the few hours of the onsite the visit or excursion itself. The task needs to be embedded within preparation and follow-up activities that contribute to the entire task. In addition to the thoughtful design of the task, the learning design must also incorporate the means to achieve the outcomes, and the preparation and selection of learning resources and supports also need to be completed. While some resources can be (and need to be) sourced by the students themselves before and after the visit, the selection of appropriate recommended resources is useful to help students commence the task activities strongly and purposefully.

In order to ensure that the design goals and learning outcomes are achievable, planned formative evaluation and user testing is important through all phases of development. Further effectiveness evaluation is also required on implementation of the learning environment. These evaluations are critical to the decision making

that occurs at various points throughout the design and development process, and later when the learning environment is used in the field. All projects described here benefited greatly from constant trialing and user testing through the development process, ensuring the design goals and learning outcomes could be achieved.

From a context for design and implementation, the design and pedagogical principles adopted for these projects need to be a whole project team understanding. This focus needs to be constantly reinforced and tested throughout the design and implementation process to ensure that projects maintains the overall vision for the implementation of authentic tasks by learners. Additionally, each project needed to clearly articulate and reinforce a set of expectations for teachers and students to ensure the pedagogical principles were constantly at the fore for all participants.

16.4 Effectively Linking the World and Classrooms

The new affordances of Web 2.0 have offered far greater innovation than has been experienced with the initial Web specifications. Young people have taken to these new affordance in their droves (ACMA 2007) through the social networking and self-publishing opportunities offered with environments and new constructs such as personal blogs, MySpace, and FaceBook. Of course, these types of experiences are not necessarily viewed by parents, educators and policy makers in a positive light with the media often focusing on the negative and legalistic aspects with sensationalizing headlines (The Sun-Herald, 28/10/07) about these settings and the challenge to authority by youth that these settings can offer. The challenge for educators here is threefold; firstly to develop an understanding of the opportunities for learning that these settings offer, secondly to educate parents and the media about the positive aspects available and thirdly to make use of these continually evolving affordances, and the enthusiasm of young people, to develop learning settings that are innovative, engaging and acceptable to educationalists, policy makers and the public. These learning settings are innovative, but not unique in that there is a growing realization by institutions that are repositories of knowledge (e.g., Museums, cf., Sumption 2006) and learning centers that technology can allow them to reach out to new and broader audiences.

16.5 Conclusion

The lessons learnt from the design and implementation process for each case reported in this chapter were often specific to each individual project, but a number of common themes around applying principles for design of authentic learning setting could be thought of as recurring. The importance of task choice, resource choice, a common sense of project goals and a robust and constant evaluation framework were common elements in these innovative environments.

All the projects described here have attempted to make use of advances in technology to construct learning settings for students that are task driven, situated and set in a social constructivist paradigm. The theoretical position taken, it is argued, will ensure high quality learning and encourage higher order thinking. The three ICT supported learning settings described all appear to be offering experiences that learners, teachers and parents all believe are exciting and challenging. The settings all make use of quite new technology affordances to effectively link learners to beyond classroom locations and experiences and offer experiences that effectively draw on situated and authentic theoretical positions.

References

- ACMA: *Media and Communication in Australian Families 2007: Report of the Media and Society Research Project* (Australian Government, Australian Communications and Media Authority, Canberra, ACT 2007).
- Agostinho, S.: Using characters in online simulated environments to guide authentic tasks. In *Authentic Learning Environments in Higher Education*, ed. by Herrington, A., Herrington, J. (Information Science Publishing, Hershey, PA 2006) pp. 88–95
- Alessi, S. M.: Learning interactive videodisc development: A case study. *J. Instr. Dev.*, **11**(2). 2–7 (1988)
- Ambron, S., Hooper, K. (Eds.): *Interactive Multimedia: Visions of Multimedia for Developers, Educators, & Information Providers*. Redmond, Washington: Microsoft Press (1988)
- Anderson, D., Lucas, K.B.: The effectiveness of orienting students to the physical features of a science museum prior to visitation. *Res Sci Educ.* **27**(4), 485–494 (1997)
- Bennett, S., Harper, B., Hedberg, J.: Designing real-life cases to support authentic design activities. In *Meeting at the Crossroads*, ed. by Kennedy, G., Keppell, M., McNaught, C., Petrovic, T., *Proceedings of the 18th Annual Conference of the Australian Society for Computers in Learning in Tertiary Education* (Biomedical Multimedia Unit, University of Melbourne. Melbourne 2001) pp. 73–81
- Bergin, D.A., Anderson, A.H., Molnar, T., Baumgartner, R., Mitchell, S., Korper, S., Curley, A., Rottmann, J.: Providing remote accessible field trips (RAFT): an evaluation study, *Comput Human Behav.* **23**(1), 192–219 (2004)
- Brickell, G., Herrington, J.: Learner challenges and situated learning: engaging students at Sydney Olympic Park. In *Beyond the Comfort Zone*, ed. by Atkinson, R., McBeath, C., Jonas-Dwyer, D. & Phillips, R., *Proceedings of the 21st ASCILITE Conference* (ASCILITE, Perth 2004), pp. 150–158
- Brickell, G., Herrington, J., Harper, B.: Technology-supported collaboration on field-based authentic tasks. In *Proceedings of EdMedia 2005 World Conference on Educational Multimedia Hypermedia and Telecommunications*, ed. by Kommers, P., Richards, G. (AACE, Norfolk, VA 2005a) pp. 1961–1968
- Brickell, G., Herrington, J., Harper, B.: Integrating problem-solving strategies and fieldwork into an authentic online learning environment. In *Proceedings of the IASTED International Conference: Education and Technology* (IASTED, Calgary, Canada 2005b) pp. 38–43
- Brickell, G. J. Herrington, J.: Understanding learning in the 21st Century: A preliminary study into mobile technologies. In *Proceedings of IADIS international conference: Mobile learning 2007*. I. A. Sanchez (Eds.), (2007), Lisbon: IADIS Press, pp. 176–179
- Brown, A.L.: Design experiments: theoretical and methodological challenges in creating complex interventions in classroom settings. *J. Learn. Sci.* **2**(2), 141–178 (1992)
- Collins, A.: Towards a design science of education. In *New Directions in Educational Technology*, ed. by Scanlon, E., O’Shea, T. (Springer, Berlin, Heidelberg, New York 1992) pp. 15–22

- Dillon, P., Prosser, D.: Educational transactions in museum online learning environments, *Int. J E-Learn.* Jan–March, 2003, 14–20 (2003)
- Din, H. W.: Using museum online resources: recommendations for a teacher training model using technology in rural Alaska. In *Proceedings of Society for Information Technology and Teacher Education International Conference*, ed. by C. Crawford et al. (AAACE, Chesapeake, VA 2005) pp. 2202–2205
- Dowling, S.: Artful connections: using virtual museum tours as a keystone for an online community of practice of K-12 art teachers. In *Proceedings of World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education*, ed. by Reeves, T., Yamashita, S. (AAACE, Chesapeake, VA 2006) pp. 497–502
- Falk, J.: The director's cut: toward an improved understanding of learning from museums. *Sci. Educ.* **88**(1), 83–96 (2004)
- Falk, J. H., Dierking, L. D.: *Learning from Museums: Visitor Experiences and the Making of Meaning* (Alta Mira Press, Walnut Creek 2000)
- Ferry, B., Kervin, L., Puglisi, S., Cambourne, B., Turbill, J., Jonassen, D., Hedberg, J.: Online classroom simulation: using a virtual classroom to support pre-service teacher thinking. In *Authentic Learning Environments in Higher Education*, ed. by Herrington, A., Herrington, J. (Information Science Publishing, Hershey, PA 2006) pp. 135–161
- Fitzsimmons, J.: Speaking snake: authentic learning and the study of literature. In *Authentic Learning Environments in Higher Education*, ed. by Herrington, A., Herrington, J. (Information Science Publishing, Hershey, PA 2006) pp. 162–171
- Gibson, S., Hooper, K.: The voyage of the MIMI. In *Interactive Multimedia, Visions of Multimedia for Developers, Educators, & Information Providers*, ed. by Ambron, S., Hooper, K. (Microsoft Press, Redmond, WA 1993) pp. 181–191
- Grant, W.C.: Wireless coyote: a computer supported field trip, *Technol. Educ. Commun. ACM* **36**(5), 57–59 (1993)
- Hein, G.E.: *Learning in the Museum* (Routledge, London, New York 1998)
- Herrington, J., Reeves, T.C., Oliver, R., Woo, Y.: Designing authentic activities in web-based courses. *J. Comput. Higher Educ.* **16**(1), 3–29 (2004)
- Herrington, J., Reeves, T. C., Oliver, R.: Immersive learning technologies: Realism and online authentic learning, *J Comput. High. Educ.*, **19**(1), pp. 65–84 (2007)
- Jonassen, D. H.: Thinking technology: towards a constructivist design model. *Educ. Technol.* **34**(3), 34–37 (1994)
- Kenderdine, S.: Inside the meta center: a cabinet of wonder. Paper presented at the Museums and the Web 1998 Conference, New Orleans, USA, 11–14 March 1999. Retrieved March 13, 2006 from http://www.archimuse.com/mw99/abstracts/prg_1062.html
- Marshall, L., Northcote, M., Lenoy, M.: Design influences in the creation of an online mathematics unit for indigenous adults. In *Meeting at the Crossroads. Proceedings of the 18th Annual Conference of the Australian Society for Computers in Learning in Tertiary Education*, ed. by Kennedy, G., Keppell, M., McNaught, C., Petrovic, T. (University of Melbourne, Melbourne 2001) pp. 113–116
- Pennell, R., Durham, M., Ozog, M., Spark, A.: Writing in context: situated learning on the web. In *What works and why: Proceedings of the 14th Annual Conference of the Australian Society for Computers in Learning in Tertiary Education*, ed. by Kevill, R., Oliver, R., Phillips, R. (Curtin University, Perth, WA 1997) pp. 463–469
- Reeves, T.C.: Socially responsible educational research. *Educ. Technol.* **40**(6), 19–28 (2000)
- Reeves, T.C., Herrington, J., Oliver, R.: A development research agenda for online collaborative learning. *Educ. Technol. Res. Dev.* **52**(4), 53–65 (2004)
- Rice, M., Owies, D., Campbell, A., Snow, R., Owen, N., Holt, D.: V-Lab: a virtual laboratory for teaching introductory concepts and methods of physical fitness and function. *Aust. J. Educ. Technol.* **15**(2), 188–206 (1999)
- Savery, J. R., Duffy, T. M. (1996). In *Constructivist learning environments: Case studies in instructional design*. B. G. Wilson (Ed.), (pp. 135–148). Englewood Cliffs, NJ: Educational Technology Publications

- Smith, P. E.: Instructional simulation: Research, theory and a case study (ED No. 267 793). (1986)
- Sumption, K.: In search of the ubiquitous museum: reflections of ten years of museums and the web, In J. Trant & D. Bearman (eds.), *Museums and the Web 2006: Proceedings* (Archives & Museum Informatics, Toronto 2006) Retrieved March 30, 2006 from <http://www.archimuse.com/mw2006/papers/>
- Terrenghi, L., Specht, M., Moritz, S.: RAFT. A role-based appliance for collaborative learning. In *Proceedings of 2AD: Second International Conference on Appliance Design* (2AD, Bristol, UK 2004)
- The Sun Herald, 28/10/07, <http://www.smh.com.au/articles/2007/10/28/1192941408684.html>

Chapter 17

Awareness and Reflection in Mobile Learning Support

Christian Glahn, Dirk Börner, Jeroen Storm, and Marcus Specht



17.1 Introduction

This chapter analyses contextual parameters for designing mobile information visualization for stimulating reflection on learning actions. We discuss the transformation process from a desktop centered web-application into a mobile web-information system. Such transformations are always required whenever desktop applications do not meet the fundamental requirements of mobile or – more precisely – hand-held technologies.

Mobile learning has been a key topic in education technology research for over a decade. Using mobile technologies for educational purposes became a major interest with the availability of consumer products, such as personal digital assistants (PDA) or cell phones. One important problem with mobile technologies in education is that these technologies are different than desktop based systems. These differences are

C. Glahn (✉)
Centre for Learning Sciences and Technologies, Open University of the Netherlands,
Valkenburgweg 177, 6419 AT Heerlen, The Netherlands
e-mail: christian.glahn@ou.nl

mainly in the areas of usability, information presentation, and computational power. Therefore, it is often not possible to transfer solutions that are well tested in desktop environments directly to a mobile platform.

With this regard, two main directions of research in the field can be identified. The first direction seeks to support instructional activities through mobile technologies (Luchini et al. 2004, 2002; Norris and Soloway 1998, 2004; Facer et al. 2004; Specht 2005; 2006). The second approach uses mobile technologies to provide contextualized learning experiences (Benford et al. 2005; Nova et al. 2005), in which the location of the users is often the key dimension for contextualization.

Our research focuses on contextualization of information presentation to support awareness and reflection in learning processes. Reflection and awareness are fundamental factors in competence development and in professionalization (Schön 1983, 1987). Both, reflection and awareness are linked to the activities of a learner and the dynamics of the learning environment. In order to support learning we choose to visualize these dynamics for the learners according to their contexts.

Previously, we designed (Glahn et al. 2007) and tested different visualizations in a desktop web-application called team.sPace (Glahn et al. 2008). These visualizations were designed and embedded in the user interface of the application as widgets. These widgets were additional information in the user interface in two ways. Firstly, the widgets contained additional information that was not accessible through the content of the application. Secondly, the widgets were displayed separately from the application's content. In other words, the additional information required additional space in the user interface.

When porting a desktop application to hand-held devices, the space of the user interface is limited. Therefore, the desktop layout has to be split into different screens. For additional information, which we call "peripheral information" splitting the interface into separate screens would mean that the information is not visible to the user by default as it was possible on the desktop. This implies that the users will have to actively access the information, through which it becomes less likely that the information is accessed at all. Therefore, alternative approaches for visualizing the dynamics of social interactions were required.

In this chapter we explore the theoretical foundations of providing graphical information on learning actions to learners, and use these foundations to understand how to embed contextual information for content into a mobile user interface. The objective is to stimulate and support awareness and reflection of mobile participants of a community about the social dynamics within their group and how help them to relate their actions to these dynamics. Mobile access offers some opportunities for reflection support that have not been tackled so far. Particular, the closeness of the mobile device to its user may help to use peer information not only in summative (post) assessment of the own actions, but also to allow formative assessment of the own actions within ongoing actions. This is directly related to Schön's (1983) concept of "reflection in action", which has not been addressed by prior research on technology enhanced learning by the means of the original author.

In the following sections, we analyze the background, concepts, and techniques that were applied in the development of the mobile group information system “team.sPod”. This system is a group information system for the iPhone and the iPod Touch, which integrates resources from web-logs and social bookmarks with conceptual and location information. First we introduce the concept of peripheral information and discuss its relevance for supporting learning through stimulating reflection; then we investigate different theories that are underlying our design of technologies for supporting reflection in collaborative learning. The section “interaction and peripheral information” we study the related research on using “interaction footprints” for supporting awareness, reflection, and participation in social software. However, we found that one of the limitations of these concepts is the notion of context, or – for that matter – its absence. In order to integrate contextual dimensions into a mobile system we analyze the notion of context in the domain of learning in the professions, mainly grounded on the concept of situated learning (Lave and Wenger 1991). Finally, we combine the different concepts and describe how they affected the integration of “peripheral information” in the design of team.sPod.

17.2 Awareness and Reflection

Both, awareness and reflection are important for learning processes and competence development. Particularly in unstructured and unguided environments, like workplaces or online communities, a person’s ability to reflect on his or her actions is an important factor for competence development (Schön 1983). Schön (1983) distinguishes two variations of reflection that are relevant to learning: reflection on action and reflection in action. In both cases the learner creates a relation between past experiences and a situation. This relation is not context free but depends among others on the learner’s goals, strategies, opinions, ideas, and history as well as on the environmental conditions in which an action takes place. The main difference of the two kinds of reflections is the time when the reflection takes place in relation to the actual action. Reflection on action refers to those activities in which the “learner” reflects on past actions. This implies that the action that is focused by the reflection has already been completed and cannot be changed by taking insights resulting from the reflection into account. In opposite, reflection in action refers to those cognitive processes that are involved in the application of knowledge and experiences for assessing and controlling an ongoing task. In this case, reflection is directly related to and intertwined with the action in progress.

The relatedness and connectedness of the reflection to the action at these levels creates different demands on the tools that can be used for supporting the learners. While with regard to reflection on action, the learners may take their time to concentrate in depth on reviewing their actions, reflection in action is part of the process and in many cases it cannot interrupt the other actions for some in depth thinking. Reflection in action is therefore part of the interactions of the learners with their

environment and the related systems. If peripheral information should be used for supporting reflection in action, it has to be integrated and embedded into the normal flow of interactions. Although interactive systems have become pervasive in working and learning environments, these systems are in many cases not designed to stimulate awareness and reflection on the situated activities of their users. Most systems model interactions as abstract workflows and rarely take the meaning of these workflows of these interactions for the users into account. In order to embed peripheral information into interactive systems, it is necessary to understand the relation of the user and the system with respect to self-organization and learning.

In that sense a single interaction is composed of two parts: an action performed by an actor and a response to this action from the system. With this regard, the use of a computer system as well as the related learning processes can be described as chains of interactions. Garries et al. (2002) define the learning interaction cycle by single interactions that are connected by the interpretation of a system's response by the learner. At this level a learning process is a flow of interactions between a learner and a corresponding system. On the level of learning interaction cycles, learning processes are considered from a micro-perspective. This model has been inspired by concepts of self-regulated learning and has many similarities with the self-regulation model found in Butler and Winne (1995). The main difference of the models provided by Butler and Winne (1995) and by Garries et al. (2002) is their focus at different aspects of motivation and engagement.

Butler and Winne (1995) ground their model on the concept of self-regulation, in which a person continuously (and consciously) evaluates or assesses his or her actions, tactics, and strategies against the responses that were received from the environment. The results of this evaluation are then used to align goals and knowledge with the current situation, in which the person is situated. The motivation for the person's follow-up action is explained through the level of achievement and satisfaction in relation to the goals and prior experiences.

Garries et al. (2002) relate their motivational model to the concept of flow (Csikszentmihalyi 1991), in which the motivation of a person is related to his or her goals as well as to complexity and difficulty of the tasks of a situation in relation to the person's perception of the situation. If a situated task is recognized as achievable and affordable depends on the person's (subconscious) estimation of the complexity and challenge of the task. The motivation for the person's follow-up actions is related to the emotional perception of a situation. This emotional perception is related to the prior experiences that are related to the current situation.

Although both models cover the reflective and the emotional nature of motivation for situated actions, they are not helpful for designing applications that take the factors into account that are highlighted by these models. This limitation is due to the scope of both models: In both cases the learning interaction cycle is limited by the focus on the learner's cognitive processes (Butler and Winne 1995; Garries et al. 2002). However, indicators for peripheral information are part of the interface of a system. In order to provide peripheral information, "the system" cannot be simplified as a black box.




Fig. 17.1 Learning interaction cycle

This limitation can be solved by including the processes of the system into the model. Following the principles of context aware systems as they are discussed in Dey & Abowd (2000), Dey (2000) and Zimmermann et al. (2005), interaction appears as a symmetrical process between an actor and a system that is interconnected by the system's interface (see Fig. 17.1). The actions of an actor on the interface are analyzed and assessed by the system. Based on this analysis the system provides a response to the actions on the interface. The actor analyses and reflects on this response to judge the results of the initial action. Further actions depend on the outcomes of this last phase (Garries et al. 2002; Beck and Wade 2004).

For learner support it is necessary to understand that the parameters that are chosen at each phase within the learning interaction cycle affect the learners' engagement and performance (Rashid et al. 2006; Garries et al. 2002, Gee 2003), and that the process is context dependent, because it is based on the learner's actions and past experiences. With regard to the learning interaction cycle, learners utilize past experiences for judging the results of their actions in the same way as peripheral information relies on the learners' interaction history.

With respect to this model the external system can stimulate reflection of a learner by raising awareness on different aspects of a situation through different responses on the user interface. Although the responses of the system are not part of the user's reflection, they can help the user to be aware about the developments and status or changes of the environment in which the user is situated. This awareness of the user is one cornerstone that fosters situated cognition and attracts users to choose appropriate steps further on in the process.

17.3 Peripheral Information

team.sPod incorporates the concept of “peripheral information” for structuring the contextualized access to the information that is available in the system. The idea of peripheral information is based on the assumption that peoples’ actions including their social participation are not entirely based on “content” related facts, but are also affected by emotions and non-functional information (Norman 2005). In the domain of educational technology the widespread interpretation is that good content and a good instructional design alone is not sufficient to stimulate learning, but it also has to be provided to the learners in an appealing and pleasant graphic and interface design. This view could lead us to the view on providing successful learning environments, which is represented by the following formula.

Successful Technology Enhanced Learning = Content + Didactic + Design.

Although this is an extremely over simplified view on the development of technology enhanced learning, this formula highlights two aspects. The first aspect is the functionalism of supporting the learning process. This means that technology enhanced learning has to be optimized with regard to learning objectives, be it those defined by a learner, those defined by an educational or training institution, or those objectives demanded for external accreditation. The second aspect addresses the success criteria for technology enhanced learning. These success criteria start playing a role, if the impact of the technology on the learning processes is not as expected. Although not reported by research, in practice it is often assumed that in these cases either one or more terms at the right hand side of the formula were not appropriate.

Not denying that content, didactics, and interface design *are* important functional dimensions that have an impact on the learning processes, the concept of peripheral information offers a different perspective that is beyond the idea of educational function. Peripheral information is independent from the given learning objectives, the contents, or the underlying didactics, but it allows learners to contextualize their activities. This information is “peripheral” with regard to the educational dimensions, which does not necessarily means that the presentation of this information is completely disconnected from the application logic. In fact, peripheral information might occur as an integral part of the content – or as it is the case in team.sPod – as the central information on the display.

In our previous research (Glahn et al. 2008) we identified that peripheral information can change the learners’ perspective on the learning activities independently from the actual content. In detail we asked a group of users to use team.sPace in two different experimental conditions. Each of the two groups received a different indicator, while the actual content and function of the portal was the same for both groups. Both indicators were based on the actions a user performed in the system. These actions were clicks on links, searches using a tag cloud, visits of the portal, and contributions to the community content. The users of the first group received an activity counter that displayed the number of their actions in a

progress bar. The users in the second group received a “performance chart” that displayed the number of their actions in relation to the best performing user as well as in relation to all other users of the system. Apart from this difference the users of both groups were able to use exactly the same functions of the system. After a period of three months using the system, the users of the first group reported on the usability of the system, while the users of the second group reported sensibly on the actions that were visible on the portal. At the same time, the users of the second group were more actively using the different functions provided by team.sSpace. Furthermore, we identified that contributing and non-contributing users of the second group interpreted the indicator differently. While contributing users were motivated and used the indicator for goal setting, non-contributing users were distracted by the indicator. Such a difference could not be observed with the first group.

This experiment indicated that it makes a difference for using a system if users receive information, which has itself no functional value to an application. At the same time this information has contextual meaning to the participants of a community. The important aspect is that peripheral information creates a relational tension in the perception of the learner, in which the visualization is translated into meaningful information.

Using these explorations it is possible to define peripheral information in contrast to the application logic of technology enhanced learning. This definition also takes the learner’s capability to control their learning process into account.

Peripheral information is any information presented alongside or is embedded in the application logic, which is independent from the objectives, from the content, or from the didactics of the ongoing learning processes, and which allows the learners to relate, to maintain, or to assess their learning activities or the activities of their peers within the current context of the learner.

The important aspect of this definition is that peripheral information is always related to the learners’ context. i.e. this information needs to allow the learners to contextualize the information and to allow meaningful interpretations that are relevant to the learners’ context. From the perspective of technology enhanced learning, peripheral information is interesting because it can be used to support awareness and stimulate reflection of learners by providing referral hooks that are independent from explicit learning objectives, contents or didactic.

17.4 Contextualization, Situation, and Learning

The concepts of group awareness indicators and social proxies that have been presented in the previous section are good examples for using peripheral information to support self-organized and self-regulated learning in complex and even unstructured environments. However, not all participants may benefit from these visualizations in the same way, as the indicators are statically related to a single context. One of

the key challenges for contextualized support for learners is to model appropriate control strategies that reflect the learners' situation and help to create meaningful relations between contents and didactics on one side, and the context in which the learners are situated. For the development of team.sPod it was helpful to relate the development to the different perspectives on context.

Lave and Wenger (1991) have introduced the terms "situated learning" and "legitimate peripheral participation". Both terms reflect the social dimension of learning. Situated learning is always embedded and contextualized by the social practices of the social community, in which the learning takes place. "Legitimate peripheral participation" refers to the process of a person becoming a fully participating and accepted member of a community through participating to this community. Therefore, learning is defined as the process of "understanding in practice" that is embedded in participating in day-to-day activities (Lave 1993, 6; Wenger 1998, 46). Lave (1993) highlights that learning is affected by cognition, action, and context, which cannot be separated from each other and equally affect the learning process. Lave and Wenger (1991) use the term context to describe both, physical and social settings in which people act.

Wenger (1998) characterizes "situated learning" as the duality of participation and reification in social practice. This relation adds two aspects to the prior work of Lave and Wenger (1991): firstly, the development of action based concepts and meaningful practices; and secondly, the notion of learning as an ongoing process, which develops through continuous negotiation of practices and concepts among practitioners. The latter aspect also indicates the limitations of the earlier concept of "legitimate peripheral participation". While "legitimate peripheral participation" is meaningful in the initial phases of learning in the context of community, the duality of practice and reification reflects the continuing learning processes of long standing members and "experts" as part of the social practice of a community (Wenger 1998, 55 pp).

Situated learning or "learning in practice" covers therefore three fundamental aspects: understanding, participation in the ongoing social activities, and context. This means that learning is a contextual problem. Intuitively, the problem can be translated into the question "what are the constitutive relationships between persons and the contexts in which they act?" This question directly addresses the relation of cognition, action, and context, which has been introduced earlier in this chapter. This relation of the three aspects is frequently used in the field of ubiquitous learning, by identifying ways of offering learning events according to the (physical) context of a learner (Wenger 1998). However, the related research on situated learning (Lave 1993) identified that the problem of context in learning has to focus on the relationships between local practices that contextualize the ways people act together, both in and across contexts. From this perspective the notion of "context" is twofold. On the one hand, context defines possible activities. On the other hand, context is defined through the activities of people. This means that learning cannot be reduced to a set of "contextual learning events", but need tight coupling to the social practices in which learning is situated.

The coupling of social practices, learning, and context is linked by Wenger (1998), (2007) to the concepts of “identity” and “meaning”. Identity refers to the concept of self, including knowledge and skills, the personal history, and the role in a social community. Meaning refers to the personal model of the world, which is used for physical and social orientation, sense making, and navigation. Both concepts are part of “socio-cultural production” (Lave 1993) and are actively constructed by the learners. This construction process is contextualized by six dimensions (Lave 1993), which we summarize as follows.

- Process
- Group or peers
- Situation and event
- Participation
- Concept
- Organization or culture (the contextual “world” of the learners)

Wenger et al. (2005) analyze the role of technology for communities of practice. This analysis mainly focused on social software that is used by online communities of practice. While this analysis focused mainly on integrated (commercial) platforms, in a newer study of the authors takes a wider perspective on social software by analyzing the use of individual tools and tool sets in terms of the Web2.0 (Wenger et al. 2005). The authors identify thirteen fundamental elements in which technology can affect the success of a community of practice – and thus can affect the learning processes within (Wenger et al. 2005, 45p). These elements have contextual functions within the collaborative learning process. These functions can be mapped onto the dimensions that were introduced earlier by Lave (1993) as it is shown in Table 17.1.

Table 17.1 Functional community elements and contextual dimensions

Lave (1993) and Wenger et al. (2005) ↓	Process	Peers	Event	Participation	Concept	World
Presence		X	X	X		
Rhythm	X		X			
Interaction		X		X		
Involvement				X		
Value		X			X	X
Connections		X			X	X
Personal identity				X		
Communal identity		X		X		
Relations		X				
Boundaries		X		X		X
Integration	X	X				
Community building	X	X		X		

This mapping of the elements identified by Wenger et al. (2005) shows that social software tools as they are built so far support the strengthening the ties between the individuals and their peers. The mapping also shows that the various elements are entirely isolated from the perspective of the context of learning. Instead, the different tools that the authors associate to these elements are sensitive to different contextual dimensions.

Given to the dualistic view on communities of practice (Wenger 1998), we identified an additional element in which technology can affect communities of practice. This element is more concept-centered than the provided elements and it can be identified in recent developments of social software, particularly in the context of the Web2.0. In Wenger's (1998) terminology this element could be called "negotiation of concepts" and refers to the value that is created for the participants of a community by negotiating on the connections of concepts to the practices that are relevant to the community. Therefore, the negotiation element relates to the contextual dimensions "peers", "participation", and "concept".

The idea of the "negotiation of concepts" is related to user generated meta-data and the bottom-up ontologies – the so called folksonomies: concepts do not exist on their own, but are social constructs. These constructs are part of the group dynamics of a community, which targets the collaborative agreement on the meaning of concepts. This is different to the value aspect, which basically aims at added values that a community generates for its participants or the surrounding world. Although the negotiation of concepts might be seen as a value for the participants, the agreement on common concepts in terms of standards, goals, or objectives is more a prerequisite for the generation of external values. It is also different from the "connections" aspect, which basically addresses the creation and maintenance of links and connections between the participants, concepts, or the external world. At the level of concepts, this aspect assumes that the participants of a community already agree on "standardized" concepts. However, as recent developments in the context of the Web2.0 show, the negotiation and standardization of concepts can be supported as a collaborative action by appropriate tools.

Examples for such tools are found in the domain of "social meta-data". The idea of social meta-data refers to the ability of a group to generate meta-data about objects and people that is more meaningful to the objectives and goals of that group than standardized taxonomies and ontologies. Particularly, tools for tagging and rating can help communities to negotiate on concepts and the relations between them as they emerge. At the same time, these tools can be used by the individual participants as meaningful resources for reasoning and reflection about the activities in the community.

17.5 Interaction Footprints and Navigation Support

Wexelblat and Maes (1999) define *interaction history* as traces of interactions between learners and objects. The authors argue that interaction history is

extensively used by learners to guide actions, to make choices, and to find things of importance or interest (Wexelblat and Maes 1999). Dron et al. (2001) use the term *footprint* to indicate the value and meaning of each interaction in creating social spaces, for which the authors introduce the term *stigmergy*. This concept has been previously applied for social navigation (Dieberger 1997). Recently, Farzan and Brusilovsky (2005) use the term *interaction footprint* to refer to different traces that are left during the interaction process. Examples for such traces are notes about accessing a document in a repository, or the time a learner spent reading a document (Farzan and Brusilovsky 2005).

As described in the previous sections indicators do not necessarily target recommendation on future actions of a learner or navigational support, but focus on meta variables of the learning process and try to foster awareness and reflection about the environment, the own learning processes or changes in them. Basically for implementing these goals a variety of contextual parameters can be important and represented in indicators. Social navigation support is a prominent example of where mainly an individual interaction footprint or status in a learning environment and footprints of peers or other users are taken as an input for indication.

17.6 Waylay and Social Affordance

Erickson and Kellogg (2003) utilize interaction footprints of chats and discussion forums for graphical presentations of social exchange in online spaces. The authors call their visualizations “social proxies”. These indicators are “minimalist graphical representations that portray socially salient aspects of an online situation”. These indicators present the status of and the relations between participants in an online environment. While doing so, a social proxy is not limited to a general view of these parameters, but also visualizes the social dynamics relative to a social space. One effect of presenting social information without implying learning activities or navigational behavior has been reported as “waylay”. “Waylay refers to the practice in which a user monitors the Cookie [a social proxy] for signs of another person’s activity [. . .], and then initiate contact” (Erickson and Kellogg 2003). The concept of waylay is different to what has been described as stigmergy (Dron et al. 2001). While stigmergy refers to pathways of activities that emerge through collaborative activities, waylay refers to virtual landmarks which are used by users to structure and plan their social activities, individually.

While “waylay” is related to a user’s observations of public spaces, Kreijns (2004) identified a similar effect for group awareness indicators on distributed activities of peer user, which the author calls “social affordance”. This effect has been observed with indicators that display the activity of other users within an online environment. Different to social proxies, these indicators are not limited to a single social space, but they provide information about the activities of users relative to the activities of their peers, without providing information how these activities are interrelated.

Social affordance refers to information that stimulates activities that are aligned to the social practice within a collaborative environment. According to the author social affordances create and depend on two relationships between the learner and the environment: the reciprocal relationship and the perception-action coupling. The reciprocal relationship is based on the social intentions of a learner and on how meaningful an environment can respond to these intentions. The perception-action coupling refers to the connection of the learners' recognitions of their environment, including the actions that they will perform in accordance to it (Kreijns 2004).

17.7 An Architecture for Peripheral Information

All approaches of using peripheral information can be mapped onto an architectural pattern. This pattern has been derived from a conceptual-architecture for context aware systems that Zimmermann et al. (2005) described. The architecture has four layers and specifies operations on the data and information flow through a system from the learner input to the system response (see Fig. 17.2). The layers are the sensor layer, the semantic layer, the control layer, and the indicator layer.

The *sensor layer* is responsible for capturing the interaction footprints. A *sensor* is a simple measuring unit for a single type of data. The objective of sensor layer is to trace learner interactions. It also includes other measures that are relevant for the learning process which are not a direct result of an interaction between the learner and the system. In the architecture the sensor layer adds data to process log in order to allow the adaptation to the interaction history.

The *semantic layer* collects the data from the sensors and from the process log and aggregates this data into higher level information. The semantic layer defines

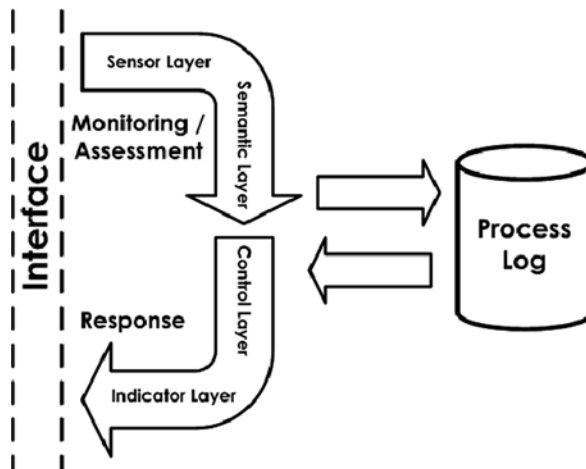


Fig. 17.2 Layers for context-aware information processing

operations or rules for processing sensor data (Cristea and Calvi 2003). A definition of how the data from one or more sensors has to be transformed is called an *aggregator* (Dey 2000). These rules are named according to their meaning, for instance *activity* or *interest*.

The aggregated information is interpreted by the *control layer* according to the history and context of a learner. The specific approach for interpretation is called a *strategy* (Cristea and Calvi 2003). It defines the conditions for selecting and combining aggregators as well as their presentation according to the learner's context. A strategy also controls the personalization of aggregators.

Finally, the aggregated information has to get presented to the learner. The *indicator layer* handles this part of the interaction. At this level the actual response is created by translating aggregated values into representations that are not just machine-readable but also accessible to humans. The active strategy of the control layer selects these representations and provides the aggregated information to them.

The first two layers are also considered as *interaction assessment* (Brusilovsky et al. 2001) or *user modeling* (Kobsa 2001). This suggests the integration of the sensor and semantic layer, although they expose different feature sets: The sensor layer is concerned with data collection of "low level information [...] including, for example, key strokes, task initiation and completion, answers of quizzes etc." (Brusilovsky et al. 2001). The main objective of the sensor layer is to organize incoming interaction footprints for further processing, while the semantic layer enriches, clusters, or transforms the data.

The last two layers are mentioned in the literature as *adaptation decision making* (Brusilovsky et al. 2001). The control and indicator layer are commonly integrated as part of the user interface. This is not always desirable because different combinations of strategies and indicators have varying effects on the learning processes and outcomes (Specht and Kobsa 1999).

This architecture is designed for context aware and context adaptive systems. When mapping the different approaches of providing peripheral information to this architecture we found that adaptation has not been considered by research, so far (Glahn et al. 2008). In these cases we can consider the control strategy as a static function that is not altered over time. Our previous research showed that this is not always optimal for user engagement and user involvement in using social software (Glahn et al. 2008b).

17.8 team.sPod

team.sPod is a mobile extension of team.sSpace (Glahn et al. 2007), particularly designed for the iPhone and the iPod Touch. team.sPod is a frontend to the team.sSpace feed aggregator that collects web-blogs and bookmarks for a group of users from the user's favorite source. All users can register their personal blog and social bookmarking source with team.sSpace and team.sPod. Therefore, the different sources of a user do not need to be stored in a central repository, but can be distributed across different services on the web. This allows the participants to choose

the communication channels that are most suitable for their needs, and then hook them into the group activities.

Given to the information provided by the participants of a group, the feed aggregator fetches the information from the different sources and makes them accessible to the participants. As each participant independently associates the personal information to a group, the other participants do not need to maintain a network of peers. Instead, the contributions of one participant become available to the entire group as soon the participant has added the references to his or her blog and social bookmarks. Additionally, the feed aggregator organizes all contributions by using the tags that are associated with them. Through a tag cloud or a tag list, users can filter the available contributions according to their interests.

Both, team.sPod and team.sPace share this basic function of the feed aggregator. The default view provided by the aggregator is a list of the most recent contributions. This basic view contains the 15 most recent bookmarks *and* the 15 most recent blog entries of the entire group. The idea underlying team.sPod was to extend this basic information provisioning with additional features that are suitable to support learning processes in mobile environments. For this purpose team.sPod integrates the localization service plazes.com in order to associate contributions not only to tags and groups, but also to locations and events. This allows user to connect social activities of blogging and bookmarking to the locations of the contributors.

Plazes.com gives users an overview about the time and location of the activities of other users. The concept behind Plazes.com consists of three core entities: *plazes*, *users* and *activities*.

- *Plazes* are named, shareable, public locations that have been marked by the users of Plazes.com. Plazes are identified by a name, a city and a country or by a hardware address of a WLAN access point. Additionally, users can tag the plazes with free form keywords.
- *Users* are the people part of Plazes.com. They are identified by their screen name, their full name and a picture.
- *Activity* describes a user's activity at a location and at a time. *Activities* are therefore the core entities at Plazes.com. At plazes.com users can create their own *Activities*, share them, and invite other participate to existing *Activities*.

As plazes.com activities connect users and locations they are used by team.sPace to identify the current location of the participants of a group.

Therefore, team.sPod is based on four general concepts: contributions, participants, tags, and locations. From a user's perspective of the iPhone or the iPod Touch, these concepts are distributed across three main screens in the system (see Fig. 17.1). This is different to the interface of team.sPace, which offers the entire information on a single screen. However, the separation of the information into different screens was not only because of screen size of the mobile devices, as one might expect. Instead, the three screens offer three different views on the social activity of the group.

The default screen a user sees is the group location view. This screen is divided in several sections, which indicate the locations where the different members of

the group are detected by *plazes.com*. This view is different from the commonly used “real world mapping” in which the absolute position of the peers is mapped on a geographically correct map. Instead, *team.sPod* locations are considered as social places and are presented to the user as different sections. Within these sections small avatars represent the people who are present at that location. The size that is devoted to a location on the screen depends on the number of persons that are currently present. The background color of a section indicates the current location of the user. The section with a light red background indicates the user’s location, while the normal background indicates otherwise. This allows the users of *team.sPod* to identify those locations that are currently important to the group. Additionally, little icons are attached to a location, if contributions were recently made by a person who is currently present at this place. From these icons a user can see immediately, which activity types are related to the persons who are present in a location.

The second screen is a list of tags. This list is ordered by the frequency this tag has been used in the system. Therefore, the most frequently used concepts are listed towards the top of the screen, while the least frequent concepts can be found towards the bottom of the screen. Tags that are applied by several participants are considered as more relevant to the group as tags that are not shared. In the list, one tag is listed per row. If the tag has been recently used with a contribution, then a small icon next to the tag indicates the type of the contribution. This allows users to identify recently used concepts, immediately. Furthermore, the background color of the line of tag indicates the user’s interest in a tag. A plain white background means that the user is not interested at all, while different colors from light blue to orange mark the different levels of interest.

Finally, the third screen is the reading view, which allows a user to read abstracts of the contributions and a link to the actual content at the original of the contribution. At this screen the user has the option to choose between the different contribution types. By default *team.sPod* offers the contribution type with the most recent entry. In this reading view only one entry is displayed at the time. By tapping at the left or right edge of the screen of the device the user can flip sequentially through the contributions.

The screens are not meant to be used in a particular order. A user may choose to flip the interface to the screen of interest. This is possible through the navigation and information bar at the top end of the screen. This bar has two buttons which refer to the two screens that are currently invisible to a user. Each button is at the far edge of the screen. Between the buttons there is an information area. This area indicates user selected filters that are currently applied to the information that is shown to the user in three rows. The last row shows the user which content type is selected for display. In the other two rows indicate if a location or a tag filter is applied, where the upper row is dominant to the lower row. This means the filters are applied in the order as they are displayed in the information bar.

A user can apply a filter by selecting a location or a tag in the related screen. In the location screen the user can tap on the section of a location on the screen in order to select that location. If a user selects a location the name associated to the location is displayed in the information bar. In the tag screen the user can select a tag by

tapping on the row of the tag. Once a location or a tag has been selected, this section is visually highlighted by a little frame. A user can unselect a selection by tapping again on the section on the screen. In this case only the related filter is removed. This means, if a tag and a location was selected, in which the tag was in dominant, and a user deselects the tag, the location filter remains. In this case the remaining filter will be in the upper row of the information bar. Alternatively, the selection can get changed by selecting a different tag or location. In this case the order of the filters remains and only the content of the filter changes.

The order of the filters in team.sPod is important, because each filter affects the content of the other screens. For example, if a user selects a location on the location screen and then switches to the tag screen, the user will see only those tags that are related to the selected location. However, if a user switches first to the tag view, selects a tag there and switches back to the location screen, the user will see only those locations in which users are present that used that tag for their contributions. This enables to users to select different perspectives to the social activity of the group.

17.9 Interaction and Peripheral Information

In the previous sections we learned that peripheral information is part of the interaction between a learner and a system, which is either a social system, such as a group of learners, or a technical system like software for computer supported training. In the past some research has focused on using content independent information for supporting the use of computer systems. Particularly, in the domains of e-commerce, computer supported collaborative work (CSCW), and computer supported collaborative learning (CSCL) information about user interactions has been used for some time already for improving the usability of a system. In the following review we will not cover the entire range of theories and solutions, rather than concentrate on the research that addressed the relation of user interactions and peripheral information. All approaches have in common that they recognize the interactions between a user and a system as a trail of events, rather than isolated statements. A system's response is therefore not only based by the current behavior of a user, but should also reflect the "interaction history".

17.10 "Situated Reflection" in team.sPod

In this last section we discuss the design concepts of team.sPod with regard to the theoretical framework that we introduced above. We choose the term "situated reflection" to indicate the relation of reflective activity and situated learning. In team.sPod peripheral information was combined with the concepts of reflection and of situated learning for supporting awareness and reflection of the system's users. The objective of team.sPod was to support both "reflection on action" and

“reflection in action” through the interface of a mobile device. The choice for utilizing mobile devices was based on the different uses of desktop computing systems and mobile systems.

These differences in using desktop and mobile devices were considered already by designing the system. The main value of mobile devices for supporting awareness and reflection is the closeness to the user. This may help to support reflection in action, when the main activity is not related to the use of a desktop computer. However, casual use and quick information are not built in functions of mobile devices but has to be considered already during the design of a system.

In team.sPod we choose to use two principles for this task. The first principle is contextualizing information to concepts. The second principle is situating the users’ location in relation to their peers. These principles are reflected in the screen showing the peer locations and the screen containing the tag list. Furthermore, we decided to minimize all features that are only of navigational purpose. This means that all screens have their own value to the users, rather than being only a step between the user and the content. Therefore we used peripheral information in the two screens in order to provide different perspectives on the group activities. A user may choose to use this information for contextualizing the information that is presented on the other screens. This helps users to create meaningful pathways to peer contributions, and allows the users to get a quick overview about the activities in the group.

17.10.1 team.sPod Support in Context

With team.sPod we focused on three of the six dimensions mentioned by Lave (1993): peers, event, and concept. Depending on the interface the dimensions are either reflected in the peripheral information or focused as the core content. Taking the stance of Wenger et al. (2005), team.sPod considers the scope of presence and connections. Furthermore, we included our additional scope of “negotiation of concepts” for the design.

When designing context-aware application one can choose between two approaches. The first approach is active and passive context awareness (Chen and Kotz 2000). Active context awareness is given if an application automatically detects a user’s context and adapts according to this context. Passive context awareness is given if an application determines the context from user inputs or selections. Passive context awareness is therefore embedded in the interaction with a user, while active context adaptation is external to the interaction between a user and a system.

For team.sPod we choose a passive approach to context awareness. This choice was made because we wanted provide the users the choice to use the tool in similar settings for varying purposes. Therefore, users define their current context by their selections in the different screens. In team.sPod contextualization is implemented with regard to two dimensions: peer location and learner interest. team.sPod determines the context based on the sequence of the selections. By default team.sPod contextualizes the content of a user according to their group associa-

tion. This means that a user has access to all contributions, locations, and tags of all members who participate in the same groups as the user. A user can choose to which dimension team.sPod should contextualize the information. If a user selects first a place, then only those tags and contributions are available, which are associated to users who are present in the selection location. However, if a user selects first a tag in the tag list then all location and contributions are contextualized regarding this selection. A user also may choose to add a sub-context that restricts the information that is available even further. In non context aware systems the information on the different screens would contain the same information regardless of the selection in one of the other screens.

For the design of team.sPod it was important that both ways of navigating through the system create meaningful pathways for the users. This means that each way of contextualizing the information allows specific interpretations of the information. Without the interpretation, both approaches would be only ways for selecting and filtering the related information. The following two use cases illustrate how this is reflected in the design of team.sPod.

The first use case focuses on the “presence” aspect of collaborative learning. In the places interface a user can see the locations that are relevant to the groups in which a user participates. The size of the location on the screen relates to the number of people who are present at the location. Additionally, users can identify their current location on the screen through its different background. This view provides an overview to the locations of a user’s peers. In this view two user questions are addressed by team.sPod: first, in which topics are the peers (currently) interested that are present at my location, and second, what are the important topics at a remote location (e.g. in the office). Both questions can be answered either by reading the contributions or by scanning the tags that were associated to the contributions.

The effect of choosing a location in team.sPod is that the information of the other screens is recalculated relative to the presence of the peers at the chosen location. This particularly means that the tag list is recalculated by taking only the present peers into account. This allows a user to identify the most important topics at a location. The additional presence of markers in the tag list for recent contributions, a user has also access to the hot topics at a location. This second feature allows the user to create connections between peers, tags and locations.

The second use case addresses the understanding and using concepts among peers. In this case a user browses the tag list, in which the most recent and the most frequently used tags are shown by default. By selecting a tag from this list, the system adapts the places screen by adding different shades of background colors that indicate the relevance of the tag at this location. Additionally, all avatars of peers that have no interest in the tag are grayed-out. This way the users can easily locate the “expertise” among their peers.

Through the visual annotations for the most recent contributions that are assigned for the most recent contributions, a user is also able to locate hot spots for a particular topic. A user can use this information to find “experts” and to relate the own knowledge to the active peers in the different locations.

17.10.2 Embedded Reflection Support

Alone, the contextualization of information will not help users to reflect in their actions. In the previous section we have already mentioned how the user interface of team.sPod changes in case of contextualization. These changes are based on peripheral information that is not part of the information provided with the resources in the information view. Earlier in this chapter we highlighted that peripheral information depends on the relations between the user and its environment. These relations are partially represented by this type of information.

The main focus for reflection support in team.sPod is based on the relation of peers, their interests and the user's own interests. This relational information is embedded in the places as well as in the tag list screen. The system keeps track of the general interests of the users by monitoring their reading and contribution habits. "Interest" is based on the occurrence of tags with the resources a user has contributed or accessed. The second aspect of team.sPod is "presence". This is only based on the current location of the user and the peers.

By combining these two aspects users become more aware about the topics that are relevant to their peers as well as to themselves. For example, user in a meeting may take the device that runs team.sPod and checks which are the relevant topics for the people who are also in the meeting. The user selects the own location as location context and switches to the tag list. Now this tag list contains only the tags that were used by the present peers. This already gives a broad overview on the interests of the group that is present in the meeting room. Based on the overlap of interests with the peers, the user can choose a better strategy for the meeting. This example also shows how team.sPod can be used for the "negotiation of concepts", as it helps to understand the relevance of topics of the different users.

At the core of reflection support is the user's awareness on different dimensions. In order to draw the user's attention to different aspects of group activities, team.sPod embeds several techniques of content annotation. This ranges from the "recently used markers" via background colors and color intensity to the size of a shape on the screen. In order to support reflection in action, it was important to present the information compact. This compactness is required by the integration into external activities, which should be minimally interrupted by the use of the mobile device. This is also reflected by the reuse of visualization concepts in the different screens. For example, the icons for recent contributions are consequently used in all screens with the same meaning for the recently uses contribution types, or the background color of a screen section always refers to the level of interest.

By embedding indicators about "presence" and "interest" in the user interface, the small screens of the iPhone and the iPod Touch are efficiently used for highlighting information to the users. At the same time the users are not forced into a single interpretation of the peripheral information that is embedded in the team.sPod interfaces. Instead, they can develop their own interpretations of what they see through the contextualization of the information, because the presented information

is independent from the underlying content of the available contributions. At the same time the peripheral information is based on the past actions of the user and the peers. This connectedness of the peripheral information with the experiences of the user supports its interpretation in the context in which the user is situated in.

17.11 Conclusion

In this chapter we discussed how the team.sPod system supports reflection for mobile learners in informal settings. The reflection support is based on two central concepts: peripheral information and learner context.

Peripheral information is information that is independent from learning objectives, but allows learners to relate their actions to the group activities. This information is “peripheral” with regard to the educational objectives and the operational application logic. However, this does not necessarily mean that the presentation of this information is completely independent from the application logic. The important aspect of peripheral information is that it creates a relational tension in the learners’ perception, which allows the learners to relate, maintain, or to assess their actions in relation to those of their peers.

The cognitive process of relating and assessing the own actions is also mentioned as reflection. We referred to Schön (1983, 1987) who distinguishes between summative reflection or “reflection on action”, and formative reflection or “reflection in action”. In both cases the learner creates a relation between an action and past experiences, knowledge, beliefs, and habits, but also to environmental conditions. Awareness about these dimensions is one cornerstone for reflection. A system can support this process by raising attention to meaningful aspects that might not be conscious to the learner. With this respect the two forms of reflection have very different demands. While learners may take their time while reflecting on action, reflection in action is embedded in ongoing activities.

From Lave and Wenger we learn that such meaning is context dependent and part of an ongoing negotiation process among peers. This leads us to one of the key challenges for supporting reflection: the modeling of appropriate interfaces that allow learners to create meaningful relations between their actions, their knowledge and experiences, and the context in which they are situated. With respect to reflection in action, this process cannot be disruptive, but has to be integrated in the normal flow of actions.

The main value of mobile devices for task integration is the closeness to the user’s body. This is particularly beneficial in those cases when the main activity is not related to the use of a desktop computer. However, task integration and non-disruptive information access are not built in functions of mobile devices, but have to be considered already during the design of a mobile system. For the design of team.sPod it was therefore important that the navigation experience creates meaningful pathways for the users.

References

- Beck, J.C. Wade M.: *Got Game, How the Gamer Generation is Reshaping the Business Forever.* (2004), Boston, MA: Harvard Business School Press
- Benford, S., Rowland, D., Flintham, M., Drozd, A., Hull, R., Reid, J., Morrison, J., Facer, K.: Life on the edge: supporting collaboration in location-based experiences, *Proc. of CHI 2005*, Portland, Oregon USA, Apr. 2005, pp. 721–730, ACM Press, ISBN 1581139985
- Brusilovsky, P., Karagiannidis, C., Sampson, D.: The benefits of layered evaluation of adaptive application and services. *8th International Conference on User Modelling*, Sonthofen, Germany (2001)
- Butler, D. L., Winne, P. H.: Feedback and self-regulated learning: a theoretical synthesis. *Rev. Educ. Res.* **65**, 245–281 (1995)
- Chen, G. Kotz, D.: A Survey of Context-Aware Mobile Computing Research. Dartmouth Computer Science Technical Report TR2000-381, (2000), Dartmouth College, Hanover, NH
- Cristea, A., Calvi, L.: The three layers of adaptation granularity. *9th International Conference on User Modelling* (Springer, Johnstown, PA 2003) pp. 4–14
- Csikszentmihaly, M.: *Flow: The Psychology of Optimal Experience* (Harper Collins, New York 1991)
- Dey, A. K., Abowd, G. D.: Towards a Better Understanding of Context and Context-Awareness. *Proc. CHI 2000 Workshop on the What, Who, Where, When, and How of Context-Awareness*, <http://ftp://ftp.cc.gatech.edu/pub/gvu/tr/1999/99-22.pdf>. (2000)
- Dey, A. K.: Understanding and using context. *Personal Ubiquitous Comput.* **5**(1), 4–7 (2001)
- Dey, A. K.: Enabling the use of context in interactive applications. *Computer-Human Interaction (CHI'00)*, de Hague, NL (2000) pp. 79–80
- Dieberger, A.: Supporting social navigation on the World Wide Web. *IJHCS*, **46**(6), 805–825 (1997)
- Dron, J., Boyne, C. W., Mitchell, R.: Footpaths in the stuff swamp. In: *WebNet 2001 World Conference on the WWW and Internet Proceedings*, 23–27 Oct. 2001, Orlando, FL, USA
- Erickson, T., Kellogg, W. A.: Social translucence: using minimalist visualizations of social activity to support collective interaction. *Designing Information Spaces: The Social Navigation Approach*, ed. by K. Höök et al. (Springer, New York, 2003) pp. 17–41
- Facer, K., Joiner, R., Stanton, D., Reid, J., Hull, R., Kirk, D.: Savannah: mobile gaming and learning? *J. Comput. Assisted Learn.* **20**(6), 399–409 (2004)
- Farzan, R., Brusilovsky, P.: Social navigation support in E-learning: what are the real footprints?, In *Proceedings of Third Workshop on Intelligent Techniques for Web Personalization (ITWP '05) at Nineteenth International Joint Conference on Artificial Intelligence*, ed. by Wade, V., Ashman, H., Smyth, B., Edinburgh, U.K., August 1 (2005) pp. 49–56
- Garries, R., Ahlers, R., Driskel, J. E.: Games, motivation, and learning: a research and practice model. *Simul. Gaming* **33**(4), 441–467 (2002)
- Gee, J.P.: What does video games have to teach us about learning and literacy. (2003), New York, NY: Palgrave Macmillan
- Glahn, C., Specht, M., Koper, R.: Smart indicators on learning interactions. In *Creating New Learning Experiences on a Global Scale: LNCS vol 4753. Second European Conference on Technology Enhanced Learning, EC-TEL*, ed. E. Duval, R. Klamma, M. Wolpers (Springer, Berlin, Heidelberg 2007) pp. 56–70
- Glahn, C. Specht, M. Koper, R.: Visualisation of interaction footprints for engagement and motivation in online communities of practice - results of first interviews. *Proceedings of the 1st Workshop on Technology Support for Self-organised Learners (TSSOLL)*, 2–3. June, 2008, Salzburg, Austria, pp. 29–43. Available at <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-349/glahn.pdf>(accessed 6 Jul. 2008)
- Kobsa, A.: Generic user modeling systems. *User Modeling User-Adapted Interact.* **11**, 49–63 (2001)

- Kreijns, K.: Socialable CSCL environments; social affordances, sociability, and social presence, doctoral dissertation, Open University of the Netherlands, Heerlen, (2004)
- Kreijns, K., Kirschner, P. A.: Group awareness widgets for enhancing social interaction in computer-supported collaborative learning environments: design and implementation. Paper presented at the 32nd ASEE/IEEE Frontiers in Education Conference, Boston, MA (2002)
- Lave, J.: The practice of learning. In *Understanding Practice, Perspectives on Activity and Context*, ed. Chaikin, S., Lave, J. (Cambridge University Press, New York 1993) pp. 3–33
- Lave, J., Wenger, E.: *Situated Learning, Legitimate Peripheral Participation* (Cambridge University Press, New York, 1991) p. 138
- Luchini, K., Quintana, C., Soloway, E.: Design guidelines for learner-centered handheld tools. *Proceedings of the Conference on Human Factors in Computing Systems*, Vienna, Austria (2004) pp. 135–142
- Luchini, K. et al.: Scaffolding in the small: designing educational supports for concept mapping on handheld computers. *Proceedings of the Conference on Human Factors in Computing systems*, Minneapolis, USA (2002)
- Norman, D. A.: *Emotional Design: Why We Love (or Hate) Everyday Things* (Basic Books, New York 2005)
- Norris, C., Soloway, E.: Using technology to address old problems in new ways. *Commun. ACM* **41**(8), 11–18 (1998)
- Norris C., Soloway E.: Envisioning the handheld-centric classroom. *J. Educ. Comput. Res.* **30**: 281–294 (2004)
- Nova, N., Girardin, F., Dillenbourg, P.: Location is not enough!: an empirical study of location-awareness in mobile collaboration. In *Proceeding of the third IEEE International Workshop on Wireless and Mobile Technologies in Educations*, Tokushima, Japan, November 28–30 (IEEE Press, New York 2005) pp. 21–28
- Rashid, Al Mamunur, Ling, Kimberly, Tassone, Regina D., Resnick, Paul, Kraut, Robert E. and Riedl, J.: Motivating participation by displaying the value of contribution. In: *Proceedings of ACM CHI 2006 Conference on Human Factors in Computing Systems* (2006) pp. 955–958
- Schön, D. A.: *The Reflective Practitioner: How Professionals Think in Action* (Maurice Temple Smith, London 1983)
- Schön, D. A.: *Educating the Reflective Practitioner* (Jossey-Bass, San Francisco 1987)
- Specht, M., Kobsa, A.: Interaction of domain expertise and interface design in adaptive educational hypermedia. *7-th International Conference on User Modeling (UM99)*, Banff, Canada (1999)
- Specht, M., Kaibel, A., Apelt, S.: *Extending LCMS for Remote Accessible Field Trips in RAFT, Third IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOMW'05)*, (2005), pp. 302–306
- Specht, M., Kravcik, M.: Authoring of Learning Objects in Context. *IJEL*, **5**(1), 25–33 (2006)
- Wenger, E.: *Communities of Practice, Learning, Meaning, and Identity*. (Cambridge University Press, Cambridge 1998)
- Wenger, E.: Learning for a small planet: community, meaning and identity. In *Open Universiteit Nederland* (Ed.), Dies Natalis (pp. 7–12). September 23, 2007. Heerlen, The Netherlands: Open University of the Netherlands. Retrieved on December 12, 2008 from <http://hdl.handle.net/1820/1077>
- Wenger, E., White, N., John, D., Rowe, K.: Technology for communities. In *Guide de mise en place et d'animation de communautés de pratique intentionnelle, CEFRIO*, Québec (2005) Also available from http://technologyforcommunities.com/CEFRIO_Book_Chapter_v_5.2.pdf (accessed 6.Jul. 2008)
- Wexelblatt, A., Maes, P.: Footprints: history-rich tools for information foraging. *CHI'99: Human Factors in Computing Systems*, Pittsburgh, PA (1999), pp. 270–277
- Zimmermann, A., Lorenz, A., Specht, M.: Personalization and Context-Management. *User Modeling User Adaptive Interact.* **15**(3–4), 275–302 (2005)

Section V

Learning Networks Integrated

Section Editor: Hubert Vogten



Section V of this book looks at infrastructure services that are needed to create and maintain an integrated Learning Network. It provides an approach to system design of an integrated architecture for Learning Networks. Readers with a technical background will benefit from these chapters by better understanding how services can be set up to enable their stakeholder professionals to utilise the tools for their professional development, either by authoring or adapting units of learning, by learning from them, or by completing an assessment.

We start off with the development of models for use cases, concepts and the domain of professional competence development. From these foundations the data model and core principles of the integrated architecture are derived.

The first chapter (Chap. 18) sets out a conceptual model of Learning Networks. Learning Networks are virtual communities of professionals, novices and training agencies in a certain professional domain, who have a shared interest to develop the competences of the professionals and novices by exchanging experience, and knowledge, by organising training events, workshops or conferences, by offering support facilities for professional development and by offering assessment facilities

H. Vogten (✉)

Centre for Learning Sciences and Technologies, Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands

e-mail: hubert.vogten@ou.nl

for prior learning, including qualifications, certificates and/or diplomas. Within this approach, four primary use cases are discussed: (a) getting a start qualification for a function or job, (b) further professional development from novice to expert, (c) how to keep up-to-date in your profession, and, (d) are there any other jobs or functions that are within my reach, given my current competences. Starting from these use cases the concepts of ‘competence’ and ‘competence development’ are discussed extensively as a model of ‘learning in context’. At the end of the chapter, a Domain Model for Learning Networks is presented that defines all the components within an architecture for Learning Networks.

The second chapter (Chap. 19) takes a detailed practical and technical view on the first release of such a Learning Network infrastructure. This infrastructure provides domain entity services that provide a backbone for the development of applications supporting Learning Networks. These services can be accessed through their APIs. In this chapter it is explained how these services can be applied. The reader is guided in detail through an example application that demonstrates all relevant facets of these services and their APIs. In some cases, the provided services might not be sufficient to meet the requirements for a particular application. Therefore, a separate section of the chapter is dedicated to the implementation of these services. An example will demonstrate how the services can be extended. After reading this chapter the reader will have ample understanding of the available domain entity services, their use and the way they can be extended with additional services.

How to Read This Section

Researchers in Information Science or related disciplines, will find Chaps. 18 and 19 of interest, where the data model and design process of the infrastructure services are described in detail from a more theoretical perspective.

If you are chiefly interested in the practical technical implementation of the data and domain model, we suggest you proceed to Section VI where examples of the implementation are described. This will be of particular interest to developers, but also to system implementers who want an overview over particular tools.

Chapter 18

A Conceptual Model of Learning Networks

Rob Koper



18.1 Introduction

In the TENCompetence project a set of UML models (Booch et al. 1999) have been developed to specify the core concepts for Learning Networks Services that support professional competence development. The three most important, high-level models are (a) the use case model, (b) the conceptual model, and (c) the domain model. The first model identifies the primary use cases we need in order to support professional competence development. The second model describes the concept of competence and competence development from a theoretical point of view. What is a competence? How does it relate to the cognitive system of an actor? How are competences developed? The third model is a UML Domain Model that defines, among other things, the components of a Learning Network, defines the concepts and relationships between the concepts in a Learning Network and provides a starting point for the design of the overall architecture for Learning Network Services, including the data model.

R. Koper (✉)
Centre for Learning Sciences and Technologies, Open University of the Netherlands,
Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands
e-mail: rob.koper@ou.nl

In this chapter we will discuss these three models with the aim to provide a deeper understanding of the core principles within the Learning Network architecture. The subsequent chapters in this book will illustrate how to use these models in concrete implementations.

18.2 The Use Case Model

It is assumed that professionals and people in general have re-occurring goals related to their employability:

- I want to study for a new type of job or a better function, e.g. my first job or because my current job doesn't offer sufficient perspectives anymore. The aim is to attain a start qualification for this new job or function ('start qualification' use case).
- I want to grow in my current job from starter to expert, or attaining better proficiency at certain competences ('job maturation' use case).
- I want to keep up to date within my current job, functions and roles, e.g., by exchanging and discussing ideas, resources and issues with other professionals ('keep up-to-date' use case).
- Is there another job within my reach, within or outside the sector I am currently employed or unemployed in ('match job profiles' use case).

These four use cases are the ones that are supported with the Learning Network infrastructure for professional competence development. Figure 18.1 provides an overview of various trajectories of learning.

In the figure, two professional sectors in society are represented (sector A and B, for example the IT sector and the Health Care sector). Each sector has a number of work tasks that should be accomplished by the professionals, some of these tasks are sector specific, others occur in more than one sector. These tasks are normally grouped into job functions that can be performed by employees, e.g. in the Health Care sector some of the tasks are assigned to nurses, some others to doctors. Furthermore, most of the functions have various function levels and variations, together representing a function group. For instance in the group 'medical doctors' you can find various levels and variants like a doctor in training (co-assistant) at the first aid or a doctor in psychiatry or a dentist. The way functions are clustered in groups is arbitrary, for instance dentists could be in a separate group.

The levels and variants of functions are related to educational levels like intermediate vocational education, professional education or university education.

Within a specific function level, the expertise of people can vary. Beginners have less expertise and are less matured in the profession than people with many years of experience. This is represented in the figure with the terms 'start level' and 'expert level'.




Fig. 18.1 Different learning trajectories related to the use cases

Within Fig. 18.1 we can now identify the use cases we discussed earlier. Getting a start qualification (use case 1) is represented by vertical movement in the figure (number 1 = initial education; number 3 = studying for a higher function level; number 6 = studying for higher qualification in a new sector, while working in another sector at a lower level).

Use case 2 (job maturation) is depicted by the number 2 in the figure. Persons are still doing the same job, but with much more experience and proficiency. Depending on a profession it can take up to six years before someone has matured in a job. Shortening this time frame is one of the promises of the use of Learning Networks.

Use case 3 (keep up-to-date) is not directly visible in the figure because it is related with a time dimension. Jobs change, environments and tools change, knowledge changes, trained skills are diminishing. When you do not keep-up with all the changes in your profession or when you stop training or practicing certain skills, then after some time, you will not be able to function anymore at the required level. So, use case 3 keeps you on a horizontal line as high as you are without falling back to lower levels of functioning. Taking care that professionals stay up-to-date is also one of the promises of Learning Networks.

Use case 4 is reflected with the numbers 4 (= other job in the same sector at the same level) and 5 (= other job in another sector at about the same level). This use case is, of course, related to employment and job mobility issues.

The core process underlying the four use cases is ‘personal development planning’ based on an assessment of past performance to estimate the competence levels of a person on the competences that are critical for a specific job or function. The

personal development plan consists of a series of formal and informal learning activities that are aimed at the fulfilment of one or more of the use cases. For instance a series of learning activities to keep up-to-date. These personal development plans can be defined strictly for personal use, or they can be embedded and agreed upon within an organisational context, like a company or school.

Within the scope of these four use cases, the conceptual model and the domain model are developed which are now discussed in the next sections of this chapter.

18.3 The Conceptual Model

The conceptual UML model makes the underlying assumptions about competence development more explicit (Fig. 18.2).

First of all it is important to understand that this model is not a technical model, although it is specified as a UML class diagram. It is only meant to serve as a conceptual model, a ‘theoretical model’ if you want.

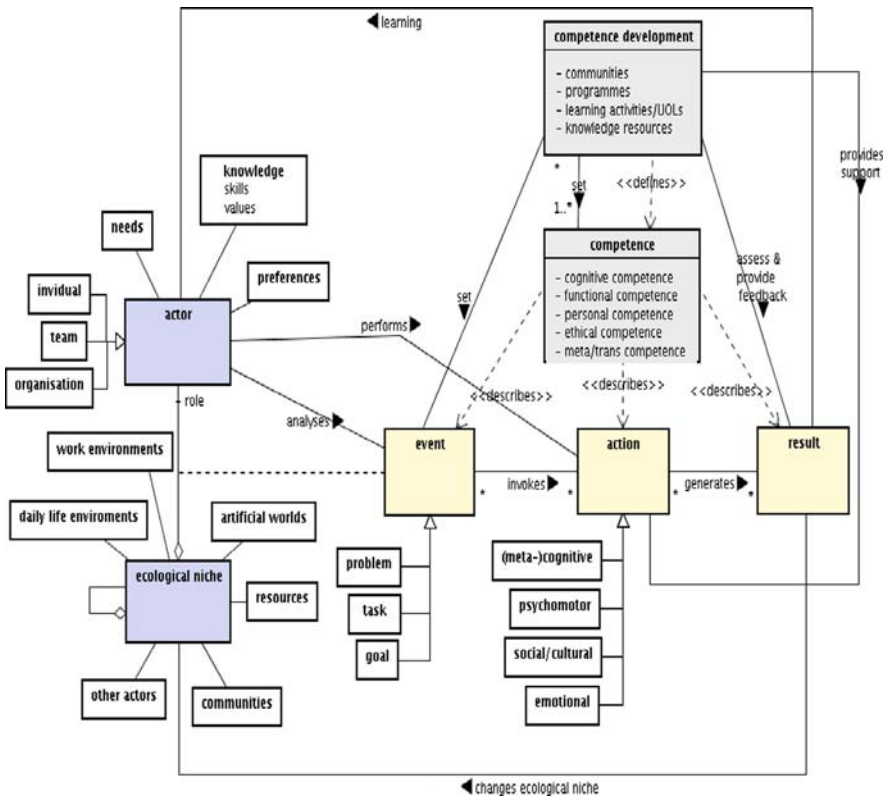


Fig. 18.2 What does a competence describe and how is competence development positioned?

The goal of the model is to specify how learning processes can be stimulated within a real life context, like work, home and school. The basic assumption underlying the model is that *'learning' occurs when a person is actively engaged in a situation in which they (re-)act on events with some observable results*. For more background reading about learning and adaptation see for instance (Schwartz and Reisberg 1991; Lieberman 2003; McInerney et al. 2004). Using this view of learning, the concepts of competence and competence development are then defined as *descriptions* of the relationship between events, actions and results.

The various components of the figure will now be explained in more detail. Let's assume there is a person, called *actor*. This actor has several (mostly non-observable) characteristics, like needs, knowledge, skills, values and personal preferences.

An actor is always *in* a so-called *ecological niche* (or 'situation' or 'context', see Hettema 1990; Odum et al. 2004). The term ecological niche is defined as: 'the specific part of the world in which an actor is at a specific moment in time, providing the context for the activities that are performed by the actor'.

Odem (1959) stated 'The ecological niche of an organism depends not only on where it lives but also on what it does. By analogy, it may be said that the habitat is the organism's "address", and the niche is its "profession", biologically speaking.' Every person, every actor has a unique set of ecological niches that they live and act in: family, home, school, work, traffic, sports, garden, study room, class room, virtual worlds (games, simulations, and dreams), etc. As has been stated before, the basic assumption that is made in the conceptual model is that learning is seen as an adaptation process: an actor acts upon the events that happen in an ecological niche, they observe and reflect upon the results and changes the actions when needed. Over time an actor learns how to act with a satisfactory result upon the different re-occurring events, and also, the actions are performed faster and more and more unconscious (also see Kolb 1984. Note: an actor doesn't have to be) an individual person; it can also represent a team or an organisation. This represents 'learning teams' and 'learning organisations'. For the rest of the chapter we will focus on the individual only. Within teams and organisations there are many more complicated layers of learning to take into account.

In most cases, actors are rather well adapted to the many ecological niches that they are involved in. Living is in this state just a kind of routine. By performing routine activities, the actions become more and more unconscious and, up to a certain limit, they are also performed faster and more reliable. This is a main effect that occurs in drill and practice training. The learning of new concepts and behaviours occurs when an actor encounters new events (tasks, problems, goals), or a completely new ecological niche. We assume that learning of new behaviours will be triggered as a result of the actions you take to deal with the event.

Events invoke *actions* in the *ecological niche* and these actions can have observable *results*. These results provide the *feedback* that is needed to learn whether the action was adequate. Note that there are many cognitive perception, attention and interpretation processes intervening with this process. For instance, although feed-

back can be observable, it does not automatically mean that a specific actor notices the feedback signals. Also the interpretation of the feedback signals can vary from person to person.

18.3.1 *The Concept of Competence*

Until now the following concepts of Fig. 18.2 have been discussed: actor, ecological niche, event, action and result. These aspects of the figure are describing real life events and the underlying assumptions of the dynamics in real life. As such it is an abstract model of the world itself: it doesn't exactly describe real experience, but it abstracts many of these types of experiences with the goal to depict the world that competences are describing. When you look at the box called 'competence' in the figure you see that this is positioned as a *descriptor* of these real world actions. So, competences are not expected to be something real within human beings, but are considered to be descriptors of the abilities of an actor to deal with some of the critical events, problems or goals that can occur in an ecological niche. For instance, being able to play tennis, being able to navigate at sea, being able to apply traffic rules, etc.

There are several questions related to competences that will now be discussed one by one:

A first question that can come to your mind is the following: *Do different people who are describing the required competences for the same job independently of each other, all come to the same (set of) competence descriptions?* The answer is of course no, first of all every model is the result of an *abstraction process* where many details are left out. Also, the level of abstraction that is used by different persons could be completely different, resulting in lists of competences that can vary enormously in length (from around 4 to hundreds of items). It is important to recognise when describing competences that they are abstractions, and the same behaviours can be described differently in terms of competences by different people. They can also be reformulated and changed at every moment in time, independent of changes in the underlying activities or situations.

This provides also an answer to the question: *How many different competences do exist?* The fact that they are descriptors means that they are part of our language. Almost all verbs in most languages are descriptors of competences: walking, singing, writing, sailing, etc. This also means that the number of competences that can be identified is about the size of the number of (contextualised) verbs that exist in a language. The contextualisation of verbs creates again new competences: walking on a rope, walking on the street, walking on the moon. Also in our language verbs are used in a more abstract or in a more concrete way. Living is a very abstract verb that includes many more detailed verbs (walking, eating, etc.). There have been various attempts to cluster competences in small groups that are independent of each other and that can be used as containers for work tasks. COLO (2009) for instance uses a core set of 25 competences to describe all core tasks in almost all of the Dutch professions at the intermediate vocational level. These are

derived from the work of Bartram (2005) who identified 8 basic competences ('the big 8'), that are expected to be responsible for the performance on work tasks. These 8 competences are:

- Leading and Deciding;
- Supporting and Co-operating;
- Interacting and Presenting;
- Analyzing and Interpreting;
- Creating and Conceptualizing;
- Organizing and Executing;
- Adapting and Coping;
- Enterprising and Performing.

In the case of the big 8 and COLO, the competences are seen as dispositions within humans (like intelligence) that are used to perform concrete work tasks. Each work task needs a combination of one or more of these competences in order to be executed effectively. The work tasks and the underlying knowledge that is needed to perform these tasks is in this approach ignored. To cope with these functional and cognitive knowledge issues, we have adapted the competence classification of Cheetham and Shivers (2005). In their work they identify several types of competences: functional, cognitive, ethical, personal and trans-/meta competences. The competences of COLO and Bartram (2005) can be seen as a combination of ethical, personal and trans-/meta competences.

Another question is: *How do you score the level of competence someone has?* The idea is that all persons have all competences up to a certain proficiency level, and that this level can be zero (representing the absence of this competence). There are many different scales on which competence proficiency levels can be measured. A standard scale has been developed in Europe, called the European Qualification Framework for professional development (EQF 2008). This framework defines eight competences levels (1–8), each representing the level of competence, the level of knowledge (cognitive competence) and the level of skill. The terminology is a bit different from the one used in this chapter. In the EQF, competence is defined as the level of autonomy you need or have to execute a task and the skill is more a definition of the complexity of the functional competence. In the concept of 'functional competence' both aspects (skill and autonomy) are combined. Table 18.1 provides an overview of the EQF levels for competence and for knowledge that can be used in conjunction with the framework we propose in this chapter. When specifying performance indicators one asks the question: 'What is the minimal level required to perform this task (at entry level and an optimal level). The difference between the entry level and the optimal level is specified as the amount of "job maturation" that is needed to perform the job as an experienced professional' (see Sect. 18.2). When assessing persons, one can ask a similar question: 'For a specific functional competence or related task, at what level does the person perform, given the evidence provided?'. When a level is assigned, one can derive that there are at least knowledge components that should be handled at the same level. This means that we do

Table 18.1 Use of the European Qualification Framework for professional development

Level	Competence	Related knowledge
0	Cannot perform the task.	No specific task related knowledge available.
1	Can perform the task under direct supervision in a structured context.	Basic general knowledge.
2	Can perform the task under supervision with some autonomy.	Basic factual knowledge of a field of work or study.
3	Can take the responsibility for completion of the task and can adapt own behaviour to circumstances in solving problems.	Knowledge of facts, principles, processes and general concepts in a field of work and study.
4	Can perform the task self-managed within the guidelines of the work context that are usually predictable, but are subject to change. Is able supervise the routine work of others, taking some responsibility for the evaluation and improvement the performance of the task.	Factual and theoretical knowledge in a broad context within a field of work and study.
5	Can perform the management and supervision in executing the task in work contexts where there is unpredictable change. The competence to review and develop performance of self and others.	Comprehensive, specialised, factual and theoretical knowledge within a field of work or study and an awareness of the boundaries of that knowledge.
6	Can perform the management and supervision in executing the task and projects at a highly complex, technical and professional level, taking responsibility for decision-making in unpredictable work contexts. The competence to take responsibility for continuing personal and group professional development.	Advanced knowledge of the field of work or study, involving a critical understanding of theories and principles.
7	Can perform the management and transformation of the task within work environments that are complex, unpredictable and require new strategic approaches. Take responsibility to contribute to the professional knowledge and practice and/or for reviewing the strategic performance of teams.	Highly specialised knowledge, some of which is at the forefront of knowledge in the field of work or study, as the basis for original thinking and/or research critical awareness of knowledge issues in a field and at the interface between fields.
8	The ability to work at the forefront in the profession, including research contexts, demonstrating substantial authority, innovation, autonomy, scholarly or professional integrity and sustained commitment to the development of new ideas or processes.	Knowledge at the most advanced frontier of a field of work or study and at the interface between fields

not measure the knowledge directly, but derive an estimated value for knowledge handling indirectly, without being specific about the field of knowledge. To be more secure on the knowledge competences, one can perform traditional knowledge tests like the ones we know from schools.

The attractive aspect of this EQF scale is that it provides a rather (life) long growth perspective on each competence, even when you have a good qualification, for instance on teaching, the framework specifies higher levels that enable you to improve. The idea is that a competence profile for a specific job contains many different functional competences that are required at various levels, e.g., for a Ph.D. Student the research competences should at the end be at level 8, but other competences, like management, teaching, etc. could be at lower levels, enabling a person to grow later in his or her career.

A next question is: *How do you specify the competences that are required in a job?* The procedure we advise to follow is the following (see Fig. 18.3):

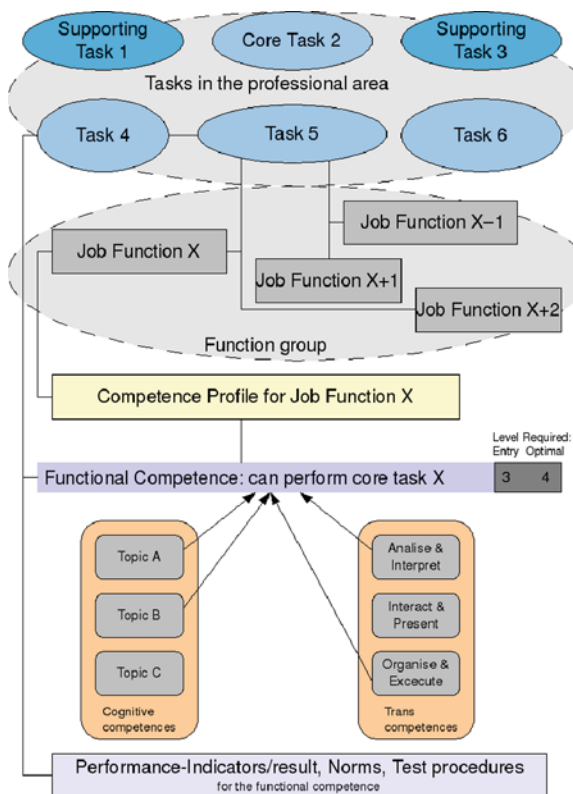


Fig. 18.3 The relationship between tasks and the various types of competences

1. First identify the core tasks in a job or in a project (e.g., write a report about the history of Holland). Core tasks may be grouped into larger categories. Besides the core tasks, also supporting tasks can be identified, like planning, communication and organisation tasks.
2. Identify the Performance Indicators to set the minimum standards for the quality and quantity of the outcome for each of the core and supporting tasks.
3. For each core and supporting task exactly one *functional competence* is assumed to be responsible for the performance of the core task (e.g., the level of your functional competence ‘report writing about the history of Holland’ is expected to be directly related to the quality level and success on the core task with about the same description). This is another way of saying that core tasks are similar to functional competences, and also the competence is measured using the performance indicators and norms for the task as an operational assessment instrument.
4. Analyse which cognitive competences (knowledge) are necessary to perform the task (history of Holland, how to structure a report, etc.). This is not an easy task, and it is also impossible to identify all the necessary knowledge requirements. It is sufficient to highlight the most significant areas of knowledge, including some of the main resources. In practice this will be implemented in a knowledge management system, e.g., by using a blogging system, a forum or any other tool that allows people to share knowledge resources, to group them into categories and to relate them to functional competences.
5. Analyse which trans-competences are needed. Cheetham and Shivers make a difference between personal, ethical, trans- and meta-competences. The demarcation between these different types of competences are not very clear, so in practice we take them together in one group that we call trans-competences (because they are not specific for the functional competences, but are used in many of them). One can use for example the big 8 or the 25 COLO trans-competences.

A further question is: *How do I know that someone has the required competences to perform a job or a task?* This question arises when hiring a new employee or when selecting the staff members for a team to do a project. When renting a sailing boat for instance, how does the rental company know that you have the required competences to sail the boat? You should provide them with some proof of your competence, e.g. by showing a license, diploma or (more informal) by telling them what ships you have sailed before. Whatever method they will be using for this assessment, they will never be completely sure that you have the competence required at a sufficient level for the circumstances you will encounter in this area, with this wind, this boat, etc. The only real test is that they let you do the job, sail the waters with their boat in various weather conditions under their observation. This is undoable of course, this is where competence assessment and accreditation comes in as a handy instrument.

Two aspects of competences are of importance here: first of all, competences are not measurable or visible in a direct way. What is measurable is the actual performance of a person, and also the recordings, the verbal or the written accounts they provide (called the ‘Performance Indicators’). The second thing is that any assessment will only *estimate* the level of competence that an actor has, it cannot provide a 100% guarantee for actual adequate behaviour.

Yet another question is: *When I attained a certain competence level, can I predict that I would behave effectively in a different situation than the one that I acquired the competence in?* This depends on the fact whether the new situation is comparable to the situation in which the competence was acquired: are there similar events (tasks, goals and problems) or do they differ quite a lot? Sailing at sea can be a completely different challenge than sailing on a river or a lake. Typically this means that some competences will be transferable and others are not: competences are situation specific. So, in any statement about a competence you should add the specific situation the competence is attained in to be able to predict future behaviour. However, the main point in the analysis of transfer of competences is whether the events that have to be tackled are similar in nature.

To conclude, the major characteristics of competences in the conceptual model are:

- A competence describes the *ability* of an actor to act effectively and efficiently when a certain class of events occur in an ecological niche.
- Competences are *abstractions* of effective and efficient actions of an actor when a certain event occurs in an ecological niche. They describe the tuple <event – action – result>. These abstractions are attributed to actors as latent variables (a kind of traits) that are predictor of the effectiveness and efficiency of the actions that an actor will perform when certain events/goals occur in an ecological niche.
- The abstraction process to create competence descriptions is similar to any *modelling process*: the same set of events, actions and results can be modelled completely different (different level of abstraction, different terminology, different selections of key events, etc.). For this reason the competences that are identified to perform a certain job well, can differ substantially depending on who is describing this job.
- The competence levels of an actor are *not observable*, but are estimated from the actual performances of actors in a real or a test situation.
- We assume that every actor has every competence to a certain *proficiency level*, including the level *zero* that represents the absence of the possibility to perform the task. In practice we use the EQF scale to specify the performance indicators and to assess the performance of persons.
- We assume that in every profession or work you can isolate some core and supporting tasks. These tasks can be clustered in larger groups and each of these tasks is related to exactly *one functional competence*. Each functional competence has a related entry proficiency level and an optimal level for experienced performers. Furthermore, each of the functional competences has some cognitive competences related to them and some trans-competences. For the trans-competences a set like the big 8 can be used.
- A competence has a *predictive value* only for comparable situations with similar types of events. An implication of this is that more abstract competences have a more general predictive use (but are harder to attain and harder to measure). Most competence profiles contain currently rather high level abstractions (about 5–12 to describe a job profile).

18.3.2 *The Concept of Competence Development*

Competence development refers to the actions of an actor in an ecological niche that intentionally or unintentionally increase his or her ability to handle certain events (tasks, problems, goals). Many activities during life are contributing to competence development in an ecological niche: handling events, learning how to handle a new technology or object, thinking, reflecting, discussing, problem solving, searching, reading, predicting, etc.

Competence development is expected to flourish well under the following conditions:

- when an actor is confronted with a problem or task within the ecological niche they are in;
- when an actor has a self-defined goal to do something new, to cope with something or to learn something;
- when an actor is confronted with a new ecological niche. In this case most events are new and are stimulating competence development by performing new activities;
- when an actor is provided with feedback and is stimulated to reflect on their actions. Specifically feedback about the (positive or negative) results of actions and improvement of adequacy or tempo of (re-)actions is of importance.

Many of the learning activities that are the result of this are informal in nature: they lead to an increase in competence, but without any prior planning, support or assessment (Livingstone 2005). Even the awareness of learning may be absent. Most of this type of learning can be characterized as *professional development* and is connected to use case 2 and 3: job maturation and keep-up-to-date. Note: a vertical jump in a career mostly requires some formal training, or at least that a person works to some minimal required competence levels for a new job or function (use case 1: attain a start qualification).

Competence development may also be the result of a planned action: a *personal development plan* (PDP): by following a course, a curriculum or programme, a training session, workshop, etc. Many of these planned events have a *learning design* that assures a serious learning process, including support. In some cases, especially for use case 1, intentional, planned learning is the best solution. In this way you can learn things that are hard to learn efficiently on your own. For instance, it can be too dangerous to put an actor in a new situation or to confront the actor with new events without any help or safety nets. In essence: efficient competence development for new function levels is stimulated when an adequate learning design is applied and sufficient support is provided.

On the other hand, many forms of professional learning (use cases 2 and 3) are more informal and develop over time without much planning or design up front. Maybe even more of our competences are developed through a combination of formal and informal learning. Also this mix can be planned with a PDP. It is exactly

this mixture of different formal and informal learning that is also the scope of the systems and methods we are developing.

18.4 The Domain Model for Learning Network Services

Our Domain Model serves several functions (cf. Fowler 2003):

- it defines the scope of the work (e.g. to build an infrastructure for professional competence development)
- it defines the vocabulary used
- it defines the relationship between the concepts used
- it defines the overall conceptual architecture
- it defines a technological theory for the project that must be tested in the pilots
- it provides a starting point for the design of other models, like the data model and web-services
- it provides the minimal functional components that must be present in an infrastructure for professional competence development

Figure 18.4 summarises the main classes of the Domain Model of a system that supports Learning Network Services (Koper 2008b). A more detailed version can be found at <http://hdl.handle.net/1820/649>.

Along with the detailed model, a description of each entity is provided. Most of these entities will be translated into concrete Learning Network Services (like a competence assessment service), including end-user tooling. In this chapter we will now shortly discuss the core concepts in the model. Please check the detailed model when you want to know more about it.

First of all the concept of a *Learning Network*, the central concept of this book. Learning Networks are virtual communities of professionals, novices and training agencies in a certain professional domain, who have a shared interest to develop the competences of the professionals and novices by exchanging experience, and knowledge, by organising training events, workshops or conferences, by offering support facilities for professional development and by offering assessment facilities for prior learning, including qualifications, certificates and/or diplomas (see also Koper and Sloep 2003; Koper and Tattersall 2004). A Learning Network is a kind of ecological niche in which events occur (mostly learning *goals* of its participants) and in which (inter-) *actions* of participants are stimulated. Most of these actions are again related to learning: following courses (called *units of learning*), doing a test, making a *personal development plan*, performing *learning activities*, searching and reading a *knowledge resource*, etc.

For all these aspects of a Learning Network tools and services are developed that are described in the subsequent chapters of this section, for instance tools and services for testing, course development and personal development planning.

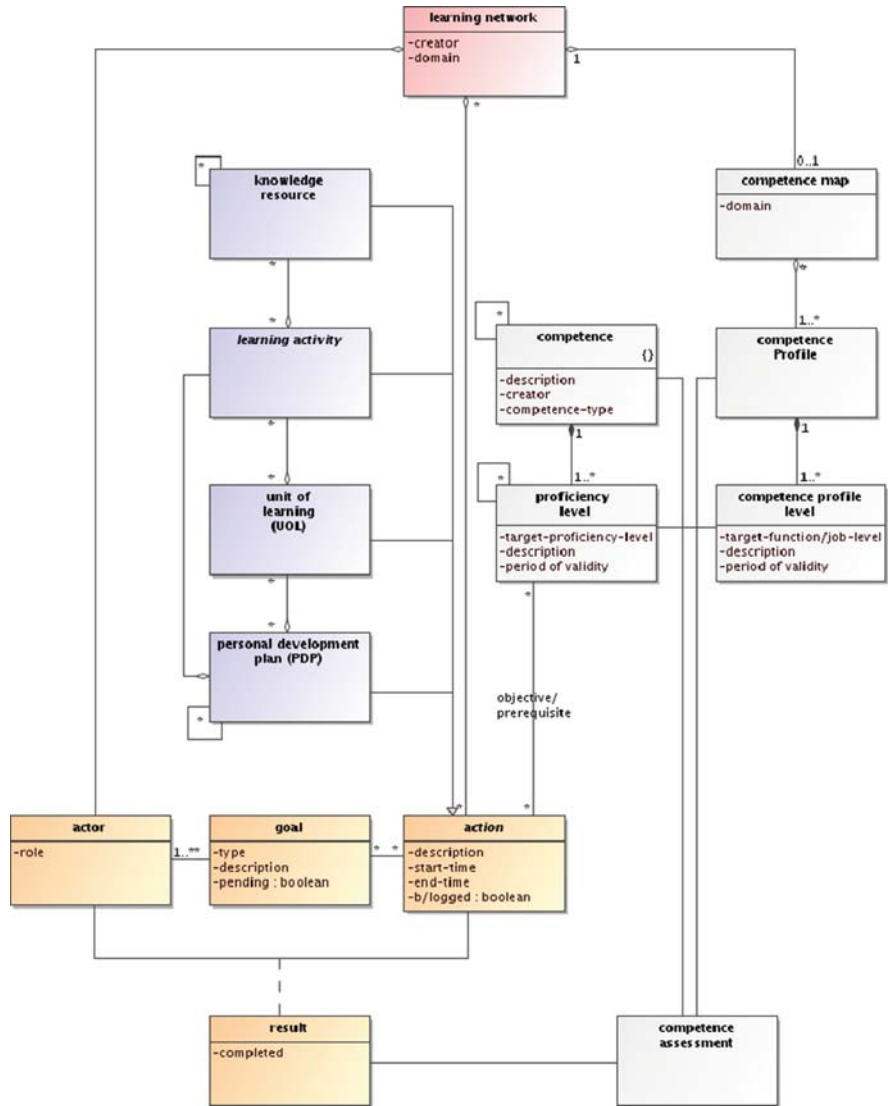


Fig. 18.4 The TENCompetence Domain Model (summary view in UML)

The (inter-)actions of the actors in a Learning Network will lead to *results* like emails, reports, ratings, wiki contributions, blogs, examination results, feedback from peers or teachers, etc. Other results are more deeply hidden in the log files of the system, like the tracks of successful learning actions (see the section on navigation in this book) that can be analysed and provided as feedback to the users in the Learning Network through mechanisms as collaborative filtering and recommender systems (Pazzani and Billsus 2007). This type of feedback stimulates what is called ‘

indirect social interactions' in the network, interactions that are based on the history of the behaviours of its participants (see e.g. Bonabeau et al. 1999; Koper 2005).

All the results, goals, assessments and actions are stored in the system to provide an e-portfolio for each user.

Technically seen, a Learning Network is a community of users who share some similar resources and services within the system. Each user is able to create their own Learning Network and can invite others to join. Given that the network is open, users can also self-enrol in a network (subscribe/unsubscribe).

Within a Learning Network a *competence map* can be created, referenced or imported. A competence map is a description of all the competence profiles for a set of professions within a domain, e.g. the professions with the domain of Aviation, or all the professions for which diplomas and degrees are registered at the Dutch government.

A *competence profile* represents a group of similar functions that can be performed at different levels. For instance in the academic sector there is the competence profile of 'academic staff' that has the levels of assistant professor, associate professor and full professor. Each function level has different requirements on the proficiency levels of the competences that are part of the profile. E.g., the minimal required competence levels to become an assistant professor at a university are different from an associate professor, but the competences are similar.

A note must be made here. Competence maps can be agreed upon at a national, regional, sector or organisation level, but there are also areas and countries where these descriptions seem to be completely absent or non-accessible. In that case some default competence maps can be assumed. In this map, each job has one job level and there is only one large competence assumed: 'be able to perform the job at that function level'. Of course this is rather non-specific, but even in this case, most of the concepts of knowledge sharing and training will still work in a Learning Network. Furthermore, in this chapter we will speak of 'professions' most of the time, although we assume that the same mechanisms that work for professions, also work in areas like hobbies, sports, personal issues, etc. that are more informally organised in society.

The visible results of the actions of a person in the network, together with the competence profiles provide the basis for the *competence assessment*. In the Learning Network system the competence assessment starts with a self-assessment of the EQF levels for a certain profile. Then a user adds evidence, e.g. based on the stored results in the e-portfolio, to prove that that level is attained. This evidence can then be used in an assessment procedure where assessors verify the level of competence. These verified competence levels are again added to the e-portfolio to form a basis for a CV that can be used to search for new jobs.

18.5 Conclusion

In this chapter, a model to support Learning Network Services for professional competence development is presented. At the basis of this work is a model that describes some highly needed use cases that should be supported in order to establish net-

work services for professional development within our society. Better tools for learning design, competence assessment, e-portfolio, informal learning, competence management, testing are needed in order to provide these services. In the next chapters the use of this model will be illustrated by explaining how tools can be build and used that are based on the model.

References

- Bartram, D.: The great eight competencies: a criterion-centric approach to validation. *J. Appl. Psychol.* **90**(6), 1185–1203 (2005)
- Bonabeau, E., Dorigo, M., Theraulaz, G.: Introduction. In *Swarm Intelligence*, ed. by Bonabeau, M. (Oxford University Press, Oxford 1999)
- Booch, G., Rumbaugh, J., Jacobson, I.: *The Unified Modeling Language User Guide* (Addison-Wesley, Reading 1999)
- Cheetham, G., Shivers, G.: *Professions, Competence and Informal Learning* (Edward Elgar Publishing, Northampton 2005)
- EQF: European Qualification Framework. http://ec.europa.eu/education/policies/educ/eqf/index_en.html (2008). Cited 27 June 2008
- COLO: Kwalificaties MBO. <http://www.kwalificatiesmbo.nl/>. Cited 27 May 2009
- Fowler, M.: *UML Distilled: A Brief Guide to the Standard Object Modeling Language*, 3rd edn. (Addison-Wesley Professional, New York 2003)
- Hettema, P. J.: *Personality and Environment: Assessment of Human Adaptation* (Wiley, New York 1990)
- Kolb, D.A.: *Experiential Learning*. New Jersey: Prentice Hall, (1984)
- Koper, E.J.R.: Increasing Learner Retention in a Simulated Learning Network Using Indirect Social Interaction. *JASSS* **8**(2), (2005). Cited 27 May 2009 at <http://jasss.soc.surrey.ac.uk/8/2/5.html>
- Koper, E. J. R.: Increasing learner retention in a simulated learning network using indirect social interaction. *J. Artif. Soc. Social Simul.* **8**(2) (2008a) <http://jasss.soc.surrey.ac.uk/8/2/5.html>. Cited 7 Oct 2008
- Koper, E.J.R., Sloep, P.B.: *Learning Networks: Learning Technology Development Programme 2003–2008*. (CELSTEC-Open Universiteit Nederland, Heerlen 2003) <http://hdl.handle.net/1820/65>. Cited 7 Oct 2008
- Koper, R.: *TENCompetence Domain Model*. (TENCompetence, Heerlen, 2008b) <http://hdl.handle.net/1820/649>. Cited 7 Oct 2008
- Koper, R., Tattersall, C.: New directions for lifelong learning using network technologies. *Br. J. Educ. Technol.* **35**(6) 689–700 (2004)
- Lieberman, D. A.: *Learning and Memory: An Integrative Approach*, 1 edn. (Wadsworth Publishing, Belmont 2003)
- Livingstone, D.W.: Basic findings of the 2004 Canadian Learning and Work Survey. Presented at the “*Future of Lifelong Learning and Work*” Conference, Toronto, June 2005. http://wall.oise.utoronto.ca/resources/WALL_Survey_Findings05.pdf Cited 7 Oct 2008
- McInerney, D. M., Alexander, P., Eiten, S. V., Fives, H., Ames, C.: *Big Theories Revisited, Research on Sociocultural Influences on Motivation and Learning*. (Information Age Publishing, Charlotte 2004)
- Odum, E. P.: *Fundamentals of Ecology* (Sanders, Michigan 1959)
- Odum, E., Brewer, R., Barrett, G. W.: *Fundamentals of Ecology*, 5th ed. (Brooks Cole, Belmont 2004)
- Pazzani, M. J., Billsus, D.: Content-based recommendation systems. In *The Adaptive Web: Methods and Strategies of Web Personalization*, Lecture Notes in Computer Science, Vol. 4321 (Springer, Berlin, Heidelberg, New York 2007) pp. 325–333
- Schwartz, B., Reisberg, D.: *Learning and Memory*, 1st edn. (Norton & Company, New York 1991)

Chapter 19

The TENCompetence Infrastructure: A Learning Network Implementation

Hubert Vogten, Harrie Martens, and Ruud Lemmers



19.1 Introduction

The TENCompetence project developed a first release of a Learning Network infrastructure to support individuals, groups and organisations in professional competence development. This infrastructure Learning Network infrastructure was released as open source to the community thereby allowing users and organisations to use and contribute to this development as they see fit. The infrastructure consists of client applications providing the user experience and server components that provide the services to these clients. These services implement the domain model (Koper 2006) by provisioning the entities of the domain model (see also Sect. 18.4) and henceforth will be referenced as domain entity services.

First, this chapter will discuss the server, its components and its intended use. The next section explains the guiding principles of the domain entity services for

H. Vogten (✉)

Centre for Learning Sciences and Technologies, Open University of the Netherlands,
Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands
e-mail: hubert.vogten@ou.nl

developing new applications for the Learning Network. The use of these services is explained in detail by taking the perspective of an example application. This example will show all concepts that are relevant when developing a client application, which uses these services. These concepts include authentication, authorisation, security, communication protocols, and XML data binding.

The final section focuses on implementation of the domain entity services. This section provides an overview of the server architecture and describes the steps involved in the processing a service call. An example explains in detail how a new service can be added to the infrastructure.

The information presented in this chapter is based on version 1.0 of the TEN-Competence Learning Network infrastructure. The corresponding API reference documentation is part of an architecture design document (Lemmers et al. 2008) and can be download at <http://hdl.handle.net/1820/1243>.

19.2 Overview Learning Network Server

A Learning Network server hosts the services for the infrastructure. The server software and a detailed installation guide can be downloaded from SourceForge at <http://sourceforge.net/projects/tencompetence>. Figure 19.1 depicts a UML implementation diagram of the server showing the major components and their relations.

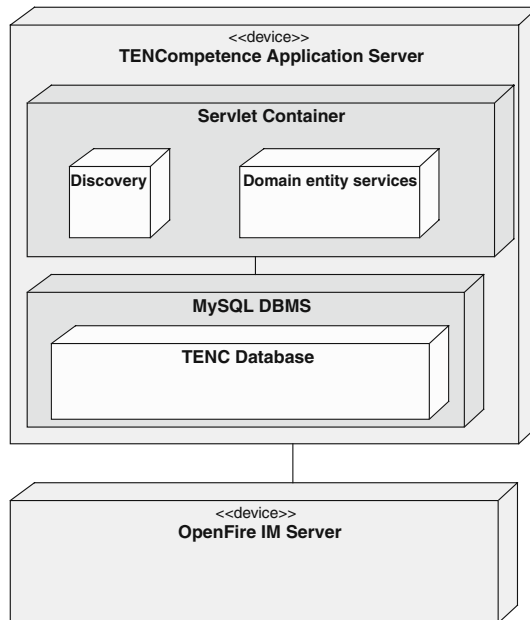


Fig. 19.1 UML implementation diagram of Learning Network server

The server uses the Tomcat application server (Apache Software Foundation 2008) for hosting a range of Java servlets Servlet. These servlets provide the entry point for the actual implementation of the services. A MySQL (2008) database is used for data storage. The server also incorporates an Openfire (2008) instant messaging server to provide chat functionality to the Learning Network infrastructure.

One of the design goals was to allow easy installation and maintenance of the server, allowing users or organisations to setup their own server and giving them control over their own infrastructure and data. As a result, a server is autonomous, and therefore needs to implement its own user management. This implies that users will need an account for each server that is hosting a Learning Network, which they want to join. An account can be shared for all Learning Networks hosted on one server.

Servers can be grouped together to form a federation. This is typically done when servers share a common topic. Federating servers allows clients to return the combined results of all federated servers for some of the queries. In the current version of the infrastructure, this federation of services is simply managed by maintaining a list of known servers on one of the server within the federation. One of the servers in this federation has to be appointed as central discovery service provider for the federation.

In the next sections we will describe how the server infrastructure can be used and extended for developing Learning Network applications.

19.3 Using the Domain Entity Services

This section is a practical guide describing how to use the domain entity services provided by the Learning Network infrastructure in a client application. We use an example application to introduce the concepts relevant when using these services. Our focus is on the use of the services, rather than on the specifics of the sample application. In the course of this example we illustrate concepts such as authentication and authorisation, discovery, and registration.

The example describes the scenario of a professional who currently is working for a recruitment centre. He is responsible for developing psychological profiles for job applicants. Although he enjoys his job, he is missing the academic challenges he experienced during his studies. Therefore, he decides to make a career change. He wants to apply for a job as researcher at his old psychology faculty. But he is not sure if he has all competences needed for the job. He has heard of a Learning Network about psychology for academics. This Learning Network maintains relevant competence profiles for this professional field and provides various learning opportunities. Therefore he wants to place himself on the researcher competence profile to find out what competences he is missing, or which of them are not up to date.

Our example application will support him in this task by allowing him to select the researcher profile from a list of all available competence profiles. For the selected profile he then can perform a self-assessment on all competences required for the profile.

The scenario mentions competences, competence profiles and assessments, which are concepts that have already been introduced in Chap. 18. The Learning Network provisions these concepts through its domain entity services Learning Network domain entity services. They are called domain entity services because they provision the elementary entities of a Learning Network. The domain entities are the common ground for all software being developed for the Learning Network; they ensure interoperability and data integration amongst additional Learning Network services and/or applications. Therefore it makes sense to reuse the services provided by Learning Network as backbone for our example application.

Before continuing our example, we describe the domain entities provided by the Learning Network infrastructure. Figure 19.2 depicts a UML class diagram of the domain entities and their relationships.

Community: community is synonymous with Learning Network in the domain model presented in Sect. 18.4; it is the central entity of the model. Users collaborate within the context of a community. All entities, except the user entity, are defined in the context of a community. This relationship is expressed by the ‘belongs to’ associations in the UML model of Fig. 19.2. This relation also has implications for the authorisation of these entities. Users that are member of a particular community have read access to all non-private entities within that community. In general, rights for entities can only be assigned to members of their containing community. Furthermore, all relations between entities are confined to the boundary of the community. So, for example, a competence defined in community A cannot be referenced in

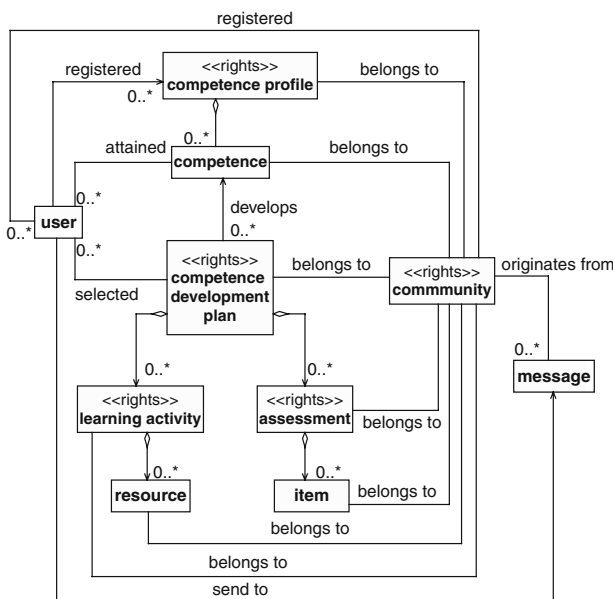


Fig. 19.2 Domain entities provided by Learning Network infrastructure

a competence profile belonging to community B. Access to a community can be controlled via the type of community and the specific access rights.

Message: users can be notified about events occurring in the system through messages. An example of such event could be the registration request of a user for a closed community. Messages are read only and they are generated by the server.

Competence profile: a competence profile is synonymous for a function or job as described in the domain model. A competence profile is defined within the scope of a community; it refers to a collection of competences, which are to be attained in order to master this competence profile. Users can register for a competence profile. The competences targeted by a competence profile can be ordered and grouped into folders, allowing meaningful organisation. The reference to a competence is attributed with a proficiency level. The proficiency levels used are informed by the European Qualification Framework (EQF) (EQF 2008) and can be represented by a numerical value ranging from one to eight. These levels are progressive and inclusive meaning that mastery of a competence for level n also implies master for level one to $n-1$.

Specific access rights can be defined for a competence profile which is represented by the 'rights' stereotype.

Competence: a competence is defined in the context of a community. The competence can be typed as personal, functional, cognitive, ethical, or trans-meta. A competence contains the attained EQF level for the user. For a competence one or more competence development plans can be defined. A competence does not contain specific rights, which implies that all community members have read and update permissions.

Competence development plan: a competence development plan is informed by the competence development programme as defined by the TENCompetence domain model. A competence development plan is defined in the context of a community and a competence; it acts as a container for a collection of learning activities and assessments that support users in attaining a particular level of a particular competence. These learning activities and assessments can be organised to form a unit of learning by structuring them into sequences and selections. A competence development plan therefore combines the competence development programme and the unit of learning from the domain model into one single entity. Many competence development plans may exist for each level of a competence. A competence development plan also defines access rights. The competence development plan also contains the progress of the user; this progress is expressed by the completion state for the plan as a whole and the completion states for the learning activities and assessments making up the plan. A user has only one active plan per competence at any moment in time.

Learning activity: a learning activity defines an activity that can be performed to achieve a particular learning goal. A learning activity is defined in the context of a community and it can belong to many competence development plans, which con-

nect the learning activity to one or more competences. Learning activities can refer to a collection of resources that are required to perform the learning activity. Access to a learning activity is defined by rights. A learning activity also contains the completion state for the user.

Resource: a resource is a container for either a reference to an external learning resource or a file that is stored on the server. It is synonymous with a knowledge resource in the TENCompetence domain model. A resource is defined in the context of a community; it has no specific rights management, implying that all community members are allowed to modify a resource. A resource does not contain any personal information.

Assessment: an assessment is synonymous with an assessment activity in the TENCompetence domain model; it refers to a collection of items. An assessment is defined in the context of a community and has rights controlling the access to this entity. An assessment includes the personal completion state for the user retrieving it. This completion indicates whether the user has taken the assessment.

Item: an item is synonymous for assessment item. Currently only multiple choice items are supported. Items are defined in the context of a community and do not contain any personal information. Items do not have specific access rights meaning that they can be read and updated by all community members.

User: the user entity represents a person on a server and contains the profile data for that person. A user can be member of multiple communities on the same server. A user entity can only be updated by the user being represented by this entity.

19.3.1 Principles of the API

Now we have seen what entities are provided by the Learning Network infrastructure we can return to our example application. One of the first things our application needs to do is to offer the user a list of all available competence profiles from which they can select the profile he wants to place himself on. Although we now know which entities to retrieve from the domain entity services, we have not yet discussed how to do this.

The domain entity services Learning Network domain entity services are implemented in accordance with the principles of REST (Fielding 2000). REST, which is an acronym for REpresentational State Transfer, is a client-server architecture style of networked systems, which outlines how resources are defined and addressed. In this architectural vision the state and functionality of an application are divided into resources. Each of these individual resources is uniquely addressable using a URI (Berners-Lee et al. 2005). A resource represents the state of an entity of the application and to modify or retrieve this entity clients communicate in a standardised way with the server to exchange representations of these resources.

Like most Restful applications, the Learning Network server uses the HTTP protocol (Fielding et al. 1999) as the underlying communication protocol. This protocol provides four methods; the GET, POST, PUT and DELETE method to determine

what action is performed on the resource. Whenever a client wants to retrieve the state of a resource it uses the GET method on the URI of the resource. We will see later how these URIs are defined in the Learning Network infrastructure. By passing the content MIME type, the client also specifies which representation of the resource it wants to retrieve. When using the 'text/xml' MIME type the server returns an xml document representing the current state of the resource. Likewise, using the content type 'text/html' would result in the server returning an html page representing the same resource. At the moment, the only representation supported by the Learning Network server is the text/xml content type.

When changing the state of an existing resource the client uses the PUT method and uploads an xml document representing the new changed state of the resource. Likewise, when the client wants to create a new resource, it uploads an xml representation of the new resource to the server using the POST method. In this case there is no existing resource whose URI can be used, so it is the responsibility of the client to generate a correct new URI that uniquely defines the new resource. Finally, when deleting a resource from the system the client requests the URI of the resource and uses the DELETE method. In this latter case no document is posted to or returned from the server.

When the client performs a REST call it must check the HTTP status code returned by the server. The server returns a status code from the 2xx range to indicate success. In case of an error it returns status codes from the 4xx or the 5xx range.

The reason for implementing the server as a collection of RESTful services is the simplicity of the stateless communication protocol as compared to a stateful protocol, which requires an additional messaging layer like SOAP (W3C 2007). This simplicity translates not only into a leaner design of the server implementation, but it also simplifies the implementation of clients using these RESTful domain entity services. An additional benefit of using a stateless protocol is the improved server scalability because there is no need to maintain session state between different server calls. The state is maintained in the data transferred between the server and the client. As there are no dependencies between two consecutive calls from a client, different servers can respond to these calls, which helps improve overall scalability.

We have now discussed the domain entities that are available and the REST protocol to access them. In the next section we will take a closer look at the details involved in using the APIs and the entities.

19.3.2 Discovery of Servers

Although by now we know the principles of the API, we still haven't seen how to retrieve the list of competence profiles. As we have discussed in Sect. 19.2, competence profiles can be hosted on different federated servers. Because servers are autonomous they can only retrieve information stored locally on the server. Therefore our application needs to collect the available competence profiles from all servers in the federation.

In order to collect the competence profiles of all servers in the federation, our application first needs to collect all the URLs of these servers. Each federation has one well known discovery service, which provides a directory of all servers in the federation. The address of this discovery service would typically be configurable in a client application. So, our application needs to retrieve the resource containing the references to all servers in the federation from this discovery service. This is achieved by performing the HTTP GET method on the URL of the discovery service resource. If the call is successful the server responds with a HTTP 200 status code and returns the XML representation of the discovery resource. Example 19.1 shows the XML document that is returned to the client. Note that for the convenience of the reader, the first lines are showing the method and relative URL used for this API call in italics, which is obviously not part of the returned XML.

```

REST method: GET
Relative URL: TENCDiscovery/servers
<?xml version="1.0" encoding="UTF-8" ?>
<servers xmlns="http://www.tencompetence.org/api/v1.0">
  <server>
    <link type="server" href="http://our.server.org/TENCServer">
      Our server
    </link>
  </server>
  <server>
    <link type="server" href="http://their.server.org/TENCServer">
      The other server
    </link>
  </server>
</servers>

```

Example 19.1 Discovery resource

The example shows that the discovery resource contains two references, one reference to ‘our server’ and one to ‘the other server’. These references are wrapped in the *link* element. Each *link* element has a *type* attribute describing the type of the resource it is referencing; a *href* attribute providing the URL of the referenced resource; and a text giving the title of the resource. This *link* element is used in many other resources as well, to create references to other domain entities.

The next step in our application is to collect the competence profiles for each server. For this purpose the server provides a resource that is listing all available competence profiles. The URL of this resource is a combination of the URL of the server plus a fixed path. The URL of the server can be retrieved from the *href* attribute of the *link* element in discovery resource and the fixed path is ‘/competence-profilesanonymous’. Applying this pattern to the first server of Example 19.1, results in the following URL: <http://our.server.org/TENCServer/competence-profilesanonymous>.

Again, fetching these resources is done via the HTTP GET method. This results in a list containing representations of all available competence profiles on the server. Example 19.2 shows the returned resource from the first server. For conciseness, the content of the second and third competence profile are omitted from the example.

```

REST method: GET
Relative URL: TENCTServer/competence-profilesanonymous

<?xml version="1.0" encoding="UTF-8" ?>
<competence-profiles xmlns="http://www.tencompetence.org/api/v1.0">
  <competence-profile>
    <link type="competence-profile" href="http://our.server.org
      /TENCTServer/com1.community/cp1.competence-profile">
      Researcher Psychology
    </link>
    <community>
      <link type="community" href="http://our.server.org
        /TENCTServer/com1.community">
        Psychology
      </link>
    </community>
    <creator>
      <link type="user"
        href="http://our.server.org/TENCTServer/user/freud.user">
        Sigmund Freud
      </link>
    </creator>
    <owner>>false</owner>
    <completed>>false</completed>
    <registered>notregistered</registered>
    <completion-count>1</completion-count>
    <registered-count>3</registered-count>
    <competence-completed-count>0</competence-completed-count>
  </competence-profile>
  <competence-profile>...</competence-profile>
  <competence-profile>...</competence-profile>
</competence-profiles>

```

Example 19.2 Competence-profilesanonymous resources

Our application collects the competence profiles of all servers in the federation and presents the resulting list to the user. From this list the user selects the competence profile they wish to place themselves on. In our example the user needs to perform a self-assessment on the set of competences of this profile. So the next step for our application is to fetch the list of competences that are part of the selected competence profile. Although this set of competences is part of the competence profile entity, the representation of the competence-profile in the list of Example 19.2 does not contain this list. Informed by REST principles, different representations of the same domain entities can be retrieved from the server. In general, two categories of representations are available.

The first category we encountered in the list of Example 19.2. This representation category is well suited for inclusion in a list, as it usually lacks details about an entity and thus forms a concise representation of the entity. This limits the amount of data transferred between the server and the client, which is especially important for lengthy lists. The other category of representation is the full representation of an entity that contains the complete data set. This full representation can be retrieved via the URL of the resource. The *href* attribute in the *link* element for an entity

contains this URL and therefore a link always points to the full representation of the entity. So in order to retrieve the full representation of an entity appearing in a list, our client application needs to simply retrieve it via the *href* from its *link* element.

19.3.3 Authentication

However, if we try to fetch the competence profile using this URL it results in an HTTP error with status code '401 Unauthorised'. As discussed in the overview, domain entities are part of a particular community and access to these entities is limited to those users that are member of that community. As long as the user hasn't joined the community, the server prohibits access to its entities. In order to grant or deny such access the server needs to establish the identity of the user doing the request and hence requires the user to pass his identity, i.e. user id, and a token to confirm this identity, the password. The Learning Network server uses the http basic authentication mechanism (Franks et al. 1999) to request these data from the user. The actual authentication is done by validating this user id and password against information from the user account stored on the server. For this process to succeed the user needs to have a user account on the server. To create a user account the client application posts an XML resource describing this new account to the server. This description contains the user id, the password and the name of the user.

Example 19.3 shows an XML description for a user resource of user 'Carl Jung'.

Because we are creating a new resource, our client application is responsible for generating the id of this resource. Informed by the REST approach taken by the Learning Network server infrastructure this new id needs to be a uniquely addressable URL. Especially the requirement this URL needs to be addressable imposes a problem; using an account name like 'c_jung' or 'carl1875' is neither unique nor addressable. And although using an email address as the user id in most cases is unique, it is not addressable.

```

REST method: POST
Relative URL: TENCServer/user/jung.user
<?xml version="1.0" encoding="UTF-8"?>
<user xmlns="http://www.tencompetence.org/api/v1.0">
  <canonical>
    <title>Carl Jung</title>
    <username>jung</username>
    <password>mypassword</password>
  </canonical>
</user>

```

Example 19.3 User resource

On the other hand using a long, unique and addressable user id can be very impractical for a user. To solve this problem the Learning Network infrastructure allows users to have a traditional short user id like the ones shown in the first example and combines this id with a server specific prefix to generate the full addressable

identifier. The pattern used to generate such a URL is <server id>/user/<user id>.user. As is the case when creating all domain entities, it is the responsibility of the client application to apply this pattern to generate the URL. Using the data of Example 19.3 this results in the following URL: `http://our.server.org/TENCServer/user/jung.user`. If our last call was successful, a new user account has been created for this user. The user id and password that were provided can now be used in all calls to the server that require authentication. Currently there are no restrictions regarding the creation of user accounts on a server.

Although the Learning Network server infrastructure allows individual servers to be federated via the discovery mechanism there is no single sign on mechanism available. Users who need to logon to different servers need to create an account on each server.

19.3.4 Authorisation

The server needs to prevent unauthorised access to its entities; it is not acceptable that each user has access to all data available on the server. For example, access to private data of other users should always be prohibited. Likewise, access to a community could be restricted to people that paid for their membership. To this effect the server implements authorisation rules. An example of such a rule we already saw in the previous section; access to domain entities is limited to users that joined the community containing those entities.

All these authorisation rules govern the different types of access allowed to the domain entities. The rules not only specify which users are allowed to read certain data, but also specify who is allowed to create or change data. There are four general rules that apply to all entities on the system:

- The first one is that all users need to be authenticated. If a user has no account on the server or if the user doesn't provide the correct password, the server prohibits all further access for this user;
- The second rule we encountered earlier; access to a community is limited to users that joined the community;
- The third rule, which is implemented via the business logic of the API, restricts data access to personal owned data only. In general, API function do not gives a user access to personal data from other users. There is one exception to this rule, which we will encounter later on. If a user requests a resource containing personal data, the server processes this request by returning the personal data from the authenticated user doing this request;
- The final rule says that the creator of a domain entity always has the right to change this entity.

Going back to our example, we created the user account but we still are not able to retrieve the selected competence profile. In order to have access to it, we first need to join the community containing this profile. Example 19.2 shows how each

competence profile references the community it belongs to through the community element. This relationship between the community and the domain entities, is emphasised by the URL of these domain entities, which takes the form <community url>/<domain-entity suffix>.

Joining a community means registering the user for the community. In order to do so, we need to change the registration state for this community. And to change the state of a resource means getting the resource from the server, modifying its state on the client and finally sending the modified state back to the server. Note that no locking mechanism is in place, meaning that all other changes made to this resource in between our get and our update calls will be lost. So first, the application fetches the community resource. The result of this call is shown in Example 19.4.

Some interesting common features can be identified from the community resource as shown in Example 19.4. First, it shows how personal state information regarding this community of the user performing the call is intertwined with canonical information about this community. The element wrappers *canonical* and *personal* are used to separate these types of information. This holds true, not only for the community domain entity, but also for most other domain entities. So in the case of our example, we can derive that the user is not yet registered for the psychology community. Another interesting feature shared by most domain entities is the *creator element*, which refers to the user entity representing the user who created this entity. This element wraps a *link* element referring to user, which therefore has attribute type set to value 'user'. By naming the wrapper element 'creator' it adds semantics to this link. Besides the canonical and personal sections, the resource also has a rights section. This section contains the authorisation rules applying to this resource.

```

REST method: GET
Relative URL: TENCServer/com1.community

<?xml version="1.0" encoding="UTF-8" ?>
<community xmlns="http://www.tencompetence.org/api/v1.0">
  <canonical>
    <title>Psychology</title>
    <description>This community...</description>
    <type>closed</type>
    <creator>
      <link type="user" href="http://TENCServer/ψ
        user/freud.user">Sigmund Freud</link>
    </creator>
  </canonical>
  <personal>
    <registered>notregistered</registered>
  </personal>
  <rights>
    <default-rights>shared</default-rights>
  </rights>
</community>

```

Example 19.4 Community resource

Earlier we saw the four authorisation rules that are imposed directly by the server and cannot be manipulated, but the rights section can be used to impose additional rules on individual domain entities. The owner of an object can specify who is granted read or update access by using these rules. For communities these rules govern whether and how users are allowed to join a community. These rules can be set or retrieved using the *rights* section from the individual resources. The *rights* section contains two optional elements. The first, *default-rights*, has the following possible values:

- ‘notshared’: the object is not accessible to others besides the creator. This in effect makes the entity hidden from view for all other users. E.g. when requesting a list of all competence profiles in a community all ‘notshared’ ones are omitted from the list. However, competence profiles created by the user requesting the list are added to the list, even they are declared ‘notshared’. Typically entities that are in their initial phase and are not deemed ready for the general public by the creator are hidden this way.
- ‘shared’: The entity can only be edited by assigned editors, or, in case of a community, access is limited and has to be approved by the appointed administrators.
- ‘sharedforupdate’: The entity can be edited by everyone, or, in case of a community, everybody is allowed to register for the community.

The second element, *update-users*, is used only in combination with the *default-rights* ‘shared’ element. In the case of a community it contains a list of all users that are appointed as the administrators of the community. For the other entities, the list contains the users that are allowed to edit. Remember that although the creator of the entity is never listed he or she is always implicitly on the list in accordance with general rule four, stating that the creator has always the right to change the entity. So, although in Example 19.4 the *update-users* list is not even present, user ‘freud’ is an administrator as he is the creator of the community. Rights can only be manipulated by the owner of the entity. It is impossible to share the right to change rights.

As can be seen from the previous description, communities are treated slightly different from the other domain entities. To emphasise this difference even further, communities have an extra data field indicating the type of access allowed to them. This field called *type* can have the following values:

- ‘private’, if a community is private, no users can join the community and the community is only visible to the creator. This is especially useful when preparing the community for deployment;
- ‘closed’, the community requires explicit approval of an administrator in order to join the community;
- and finally ‘open’ meaning any user may register for this community.

Returning to our example application; in order to register and join the community our application puts the changed state back to the server after changing the value for

the *registered* element in the personal section of the resource from ‘notregistered’ to ‘registered’. Example 19.5 shows the resource being put to the server.

```

REST method: PUT
Relative URL: TENCServer/com1.community

<?xml version="1.0" encoding="UTF-8" ?>
<community xmlns="http://www.tencompetence.org/api/v1.0">
  <personal>
    <registered>registered</registered>
  </personal>
</community>

```

Example 19.5 Community resource as put to the server

As the server uses the information provided during the authentication to determine which user is trying to register, the user identification does not have to be part of the resource XML. Also note that only the changed *registered* element and its wrapper were sent to the server. The server only allows updating those parts of a domain entity for which the user has update authorisation. In our case, this applies only to the elements in the personal section. Details about which elements can be updated are specified in the API (Lemmers et al. 2008). It is best practice to send only changed elements to the server because this improves performance on the server. Some elements, such as the *creator*, are read only and cannot be modified after initial creation. These elements are not allowed to be included in a PUT operation and the server aborts the update operation if these elements are present.

The server processes the registration request as sent in Example 19.5. However, the server must adhere to the business rules of the community. In our Example 19.4, the community *type* element was set to ‘closed’. So, the registration request cannot be handled immediately but requires human intervention from an administrator of the community. The server will set the registered state for the user to ‘pending’ to indicate that the user has not yet joined the community and is still awaiting approval. The community resource with this updated state is returned as response to the client’s PUT operation. The administrators of the community can either grant or deny user access to the community; the registered state for the community will become either ‘registered’ or ‘rejected’. It is the responsibility of our application to regularly check the registration state for the community by fetching the community resource. Example 19.6 shows the community resource returned by the server as a result of the registration request.

Besides responding with this pending state, the server also prepares a message for the administrators of the community. As we saw earlier, these administrators include both the creator of the community, as well as all users having update-rights. The message notifies the administrators of Carl Jung’s attempt to join the community. Messages can either be posted to individual users or they can be posted to all users

```

REST method: GET
Relative URL: TENCServer/com1.community

<?xml version="1.0" encoding="UTF-8" ?>
<community xmlns="http://www.tencompetence.org/api/v1.0">
  <canonical>
    ...
  </canonical>
  <personal>
    <registered>pending</registered>
  </personal>
  <rights>
    ...
  </rights>
</community>

```

Example 19.6 Pending state is returned for user Carl Jung

in a particular community. An example of the latter is the message that is posted when a new user joins the community. When addressing all users of a community the server does not post separate messages to all individual users in the community but instead addresses the message directly at the community using the id of the community as the recipient. All posted messages are kept on the server and must be retrieved by the client.

Messages can be retrieved using a GET method from the message service. The URL is `http://our.server.org/TENCServer/messages`. This will return a list of all messages sent to this user. When the application wants to retrieve the messages addressed directly to a community it adds the *recipient* parameter to the call specifying the id of the community. In this case the server returns a list of all messages sent to the community. These lists can be extensive but can be filtered via an additional parameter. The parameter *offset* filters the messages on the basis of their numerical message id; the offset parameter acts as minimum threshold number. Because messages are incrementally numbered, it can be useful to use the id of the last fetched message as offset parameter. This way, only newer messages are returned. It is the responsibility of an application to check for new messages at regular intervals. Adding both parameters to our URL would result in the following URL: `http://our.server.org/TENCServer/messages?offset=1&recipient=http://our.server.org/TENCServer/com1` In our example, Sigmund Freud has launched the application and the data shown in Example 19.7 are returned when the application fetches the messages.

This notifies user Sigmund Freud about Carl Jung's request to join the community. He decides to grant access and our application reacts by setting the registered state for Carl Jung to 'registered' for community 'com1'. Note that this is one of the very few occasions where a user sets personal data for another user. Setting the registration state is achieved via the POST method using the

```

REST method: GET
Relative URL: TENCServer/messages?offset=1
              &recipient=http://our.server.org/TENCServer/com1

<?xml version="1.0" encoding="UTF-8" ?>
<messages xmlns="http://www.tencompetence.org/api/v1.0">
  <registration-request id="2">
    <requester>
      <link type="user" href="http://our.server.org←
        /TENCServer/user/jung.user">Carl Jung</link>
    </requester>
    <community>
      <link type="community" href="http://our.server.org←
        /TENCServer/com1.community">Psychology</link>
    </community>
    <date>2008-06-13T11:40:09.000+02:00</date>
  </registration-request>
</messages>

```

Example 19.7 Messages retrieved from server for community ‘com1’

following URL: <http://our.server.org/TENCServer/com1.community/registration>. Example 19.8 shows the data that has been posted.

If Carl Jung would have been denied access to the community, the registration state would have simply been set to ‘rejected’. As a result, the server sends the outcome of this registration, be it acceptance or rejection, as a message to Carl Jung.

```

REST method: POST
Relative URL: TENCServer/com1.community/registration

<?xml version="1.0" encoding="UTF-8" ?>
<registration xmlns="http://www.tencompetence.org/api/v1.0">
  <requester><link type="user" href="http://our.server.org←
    /TENCServer/user/jung.user"/>
  </requester>
  <registered>registered</registered>
</registration>

```

Example 19.8 Registration data posted to let Carl Jung join community ‘com1’

As we saw in Example 19.4 only Sigmund Freud, who created the Psychology community, is an administrator of this community. He might decide to share his administrative privileges with Carl Jung, how unlikely this may seem at first glance. For this purpose our application will provide a separate dialog allowing the creator to add administrators to the community. In order to retrieve the data for this dialog, our application performs a GET using one of the list services provided by the server. A list service typically returns a collection of entities and acts as entry point to discover available entities. To list all users of the community our application uses the following URL: <http://our.server.org/TENCServer/com1.community/users>. The output of this call is shown in Example 19.9.

```

REST method: GET
Relative URL: TENCServer/com1.community/users

<?xml version="1.0" encoding="UTF-8" ?>
<users xmlns="http://www.tencompetence.org/api/v1.0">
  <user>
    <link type="user" ref="http://our.server.org
      /TENCServer/user/freud.user">Sigmund Freud</link>
  </user>
  <user>
    <link type="user" href="http:// our.server.org
      /TENCServer/user/jung.user">Carl Jung</link>
  </user>
</users>

```

Example 19.9 User list for psychology community ‘com1’

After Carl Jung has been selected from the list, our application must modify the rights of the community. User ‘Carl Jung’ is added to the persons having the right to control who can join the community. For this purpose a reference to Carl Jung is added to the users with update rights. Example 19.10 shows the changed community resource that is PUT to the server.

```

REST method: GET
Relative URL: TENCServer/com1.community

<?xml version="1.0" encoding="UTF-8" ?>
<community xmlns="http://www.tencompetence.org/api/v1.0">
  <rights>
    <default-rights>shared</default-rights>
    <update-users>
      <editor>
        <link type="user" href="http:// our.server.org
          /TENCServer/user/jung.user">Carl Jung</link>
      </editor>
    </update-users>
  </rights>
</community>

```

Example 19.10 ‘Carl Jung’ now also controls who can join this community

19.3.5 Putting It All Together

As soon as the registration process is finished successfully, the application is able to fetch the competence profile by providing the credentials for the user when performing the HTTP GET method. Example 19.11 shows the results of this call.

Next, our application will register the user for this competence profile. This is done in a similar fashion to the registration for the community, with the exception

```

REST method: GET
Relative URL: /TENCServer/com1.community/cpl.competence-profile

<?xml version="1.0" encoding="UTF-8" ?>
<competence-profile xmlns="http://www.tencompetence.org/api/v1.0">
  <canonical>
    <title>Researcher Psychology</title>
    <description>To be appointed an associate prof...</description>
    <community>
      <link type="community" href="http://our.server.org/␣
        TENCServer/com1.community">Psychology</link>
    </community>
    <creator>
      <link type="user" href="http://our.server.org/␣
        TENCServer/user/freud.user">Sigmund Freud</link>
    </creator>
    <ordered-competence-list>
      <competence-level>
        <level>6</level>
        <competence>
          <link type="competence" href="http://our.server.org/␣
            TENCServer/com1.community/c1.competence">
            Project Management
          </link>
        </competence>
      </competence-level>
      <competence-level>
        <level>5</level>
        <competence>
          <link type="competence" href="http://our.server.org/␣
            TENCServer/com1.community/c2.competence">
            Research funds raising
          </link>
        </competence>
      </competence-level>
    </ordered-competence-list>
  </canonical>
  <personal>
    <completed>>false</completed>
    <registered>notregistered</registered>
  </personal>
  <rights>
    <default-rights>sharedforupdate</default-rights>
  </rights>
</competence-profile>

```

Example 19.11 Competence profile resource

that a server will never deny this registration. So the only valid states for the *registered* element are ‘registered’ and ‘notregistered’. This holds true for all other domain entities for which a user can be registered with the exception of the community entity. We omit further details concerning this call, because the process is very similar to the community registration of Example 19.5.

Our application now has access to the set of competences that are part of this competence profile. They are ordered by the *ordered-competence-list* element,

which reflects the ordering of these competences by the author of this competence profile. Each competence is wrapped in a *competence-level* element, which associates each competence with a target level. Such a target level represents the level that has to be attained for the competence to fulfil the requirements of the competence profile. The levels are informed by the EQF and the eight levels are represented by their numerical values (see also Sect. 18.3.1). Example 19.11 shows two competences of which the first needs to be mastered at level 6 and the second at level 5.

Our application will show each individual competence and the corresponding required levels. In order to show additional information about the competences such as their descriptions and also the level the user has attained, the application retrieves these competences individually. An example of a competence resource is shown in Example 19.12.

```

REST method: GET
Relative URL: /TENServer/ com1.community/c1.competence

<?xml version="1.0" encoding="UTF-8" ?>
<competence xmlns="http://www.tencompetence.org/api/v1.0">
  <canonical>
    <title>Research fund raising</title>
    <description>Research funds are essential...</description>
    <type>functional</type>
    <community>
      <link type="community" href="http://our.server.org/
TENServer/com1.community">Psychology</link>
    </community>
    <creator>
      <link type="user" href="http://our.server.org/
TENServer/user/freud.user">Sigmund Freud</link>
    </creator>
  </canonical>
  <personal>
    <owner>>false</owner>
  </personal>
</competence>

```

Example 19.12 Competence resource

The level a user has attained for a competence is returned in the *personal* section of the competence resource in the *attained-level* element. This level is represented by the numerical value for the corresponding EQF level. Note that Example 19.12 does not contain the *attained-level* element. This implies that a server has no state information regarding the attained level for this user. In principle empty values or null values are omitted by the server. This applies to all domain entities. Reason for this behaviour is the possibility to distinguish an empty string from a null value.

Seeing that they have no attained level for the competence, the user decides to assess which level of the competence he has attained. This assessment process is out of scope for this example and will not be elaborated here. But once this process is completed the user knows at which level they attained the competence and enters

this level in the example application. Next our application will send the updated competence resource back to the server. Example 19.13 shows the resource being put to the server.

```

REST method: PUT
Relative URL: /TENCServer/ com1.community/c1.competence

<?xml version="1.0" encoding="UTF-8" ?>
<competence xmlns="http://www.tencompetence.org/api/v1.0">
  <personal>
    <attained-level>3</attained-level>
  </personal>
</competence>

```

Example 19.13 Updated competence resource

This final step completes the example. Although this example application is a far stretch from a real world application, it did show all relevant concepts that come into play when developing a client using the domain entity services.

19.4 Inside the Domain Entity Services

In the previous sections we described in great detail how to use the domain entity services. However, there might be situations in which the services provided by the Learning Network infrastructure are not sufficient to build a particular application. Additional services or changes to existing services may be required. For example, there could be a need for a service that is capable of matching competence profiles with job descriptions as specified by employment agencies. The Learning Network infrastructure has been developed to allow new services to be added. This section will discuss the implementation of the domain entity services and provide an example explaining how such a new service can be added to the infrastructure.

The implementation of the domain entity services is discussed in two parts, each providing its own view on the system. The first part takes us through the main process of the services – handling messages – and introduces the system’s most important classes. The second part shows how the system’s classes are organised by providing a UML package diagram and explaining the contents of each package.

19.4.1 The Main Process: Handling Messages

The implementation of the domain entity services forms the subsystem handling the requests from the different client applications. Based on the REST principles, it offers create, read, update and delete services for the different domain entities via exchanging XML messages. Handling these XML messages is the core process for

all domain entity services. When receiving a message, the system processes it and returns a response message. The exchanged messages are all in an XML format; it is up to the client to present the received response message to users.

Figure 19.3 is a graphical representation of the message handling process and the main classes involved. Let’s look at each of the steps in the process to get a better understanding on how the system works, which will be essential when extending the services.

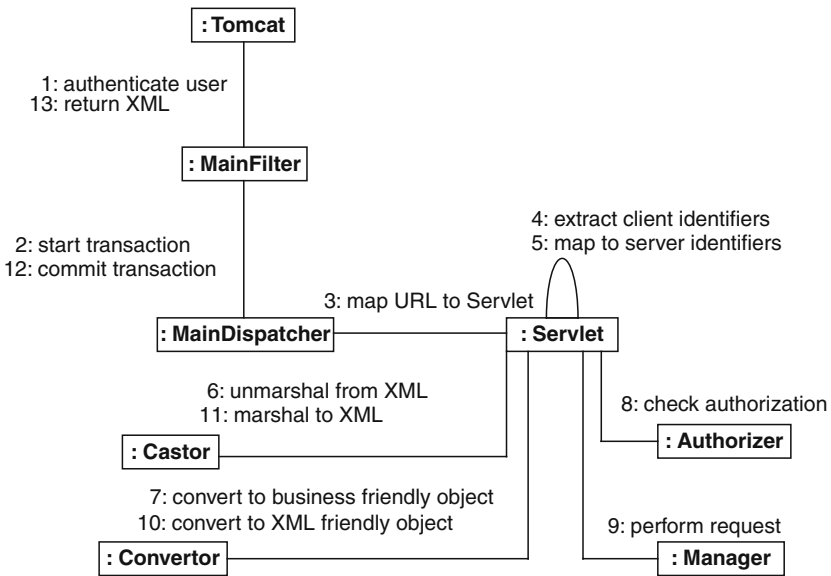


Fig. 19.3 UML activity diagram of the message handling process

When a client sends a request to the server, the very first step is a check of the user’s authentication (step 1) by the servlet container. Only after successful authentication is the request sent to the domain entity services subsystem.

The *MainFilter* filter picks up the request, starts a new database transaction (step 2) for this request and determines the logged on user before passing it to the *MainDispatcher* servlet.

The *MainDispatcher* servlet serves two purposes. The first is mapping URLs to servlets (step 3) using regular expressions, because the standard functionality offered for this mapping by the Tomcat servlet container is too limited for the domain entity services naming scheme. The second is extracting the identifiers (step 4) embedded in the request URL. After mapping the URL to a target servlet, the *MainDispatcher* passes control to the found target servlet.

The found target servlet translates the client identifiers extracted by the *MainDispatcher* to an internal format (step 5). The identifiers in the URL are unique identifiers, often implemented as GUIDs (Leach et al. 2005), generated by the different clients, whereas the server internally identifies an object by combining an object

type and an integer identifier (e.g. the object type determines a database table and the integer the object's id in that table) and thus a lookup needs to be performed. It then uses the Castor (2008) library to unmarshal the XML from the request message to a Java object (step 6). A configuration file describes to Castor how to map each type of XML message to a Java class.

The thus instantiated Java object very closely resembles the format of the XML message. The format needed by the business logic differs in a number of ways. To accommodate easy and presentation independent handling in the business logic, the instantiated Java object is converted to a business logic friendly object (step 7) by a *Convertor* helper class, which uses the Dozer (2008) library for this conversion. Dozer is a bean mapper, mapping objects from one class onto objects of another class, based on naming similarities and the contents of a configuration file. The bean mapper prevents writing and maintaining many small and very similar conversion routines yourself.

After conversion, the servlet passes control to a business logic class. The exact class depends on the type of object(s) that will be manipulated. Each type of object has its own *Manager* class, containing the business logic code for that type of object. The *Manager* class needs to do one more thing before the "actual work" can finally start. The authentication in step (1) assured the logged on person is a user of the system, but it does not assure he/she is also allowed to perform the requested action for the specified object. The *Manager* uses the *Authorizer* class to check if the requested action on the specified object is allowed by the current user (step 8) to prevent actions like inadvertent changes to data or retrieval of other users' private data. After authorisation, the *Manager* class performs the actions that are required to fulfil the client's request (step 9). These actions range from a simple retrieval of a single database row to a complex calculation. Any change to the database is handled by the Hibernate (2008) library, which offers extensive functionality for object-relational mappings.

Before marshalling the retrieved results from Java objects to XML (step 11) by the Castor library, the results from the business logic call is converted to a format closely matching the XML structure of the response message (step 10) by one of the *Convertor* helper classes.

The action requested by the user has been executed by the domain entity services and control is returned to *MainFilter* which commits the transaction (step 12) making any database changes permanent.

When a functional or technical error occurs in any of the previous process steps, the server will return an error message to the client and use one of the standard HTTP status codes to indicate the type of error. The 4xx range is used for functional errors and the 5xx range for technical errors. For example, if no object exists for the URL sent by a client, the server will return the HTTP 404 "Not Found" status code. When no errors occurred in any of the previous steps, the server will return an HTTP 200 "OK" status code and return the XML (step 13) as the body of the HTTP response, thereby completing the process.

19.4.2 Grouping: Logical Parts of the System

Besides understanding the process flow, we also need to know where to find the different parts of the code. Figure 19.4 shows the domain entity services subsystem in a package diagram, conveying the system’s different layers and how the system’s main Java packages are spread over these layers. Table 19.1 provides an additional level of detail by describing the contents of each package. These descriptions also map the steps and classes mentioned in the message handling process to the packages containing their code.

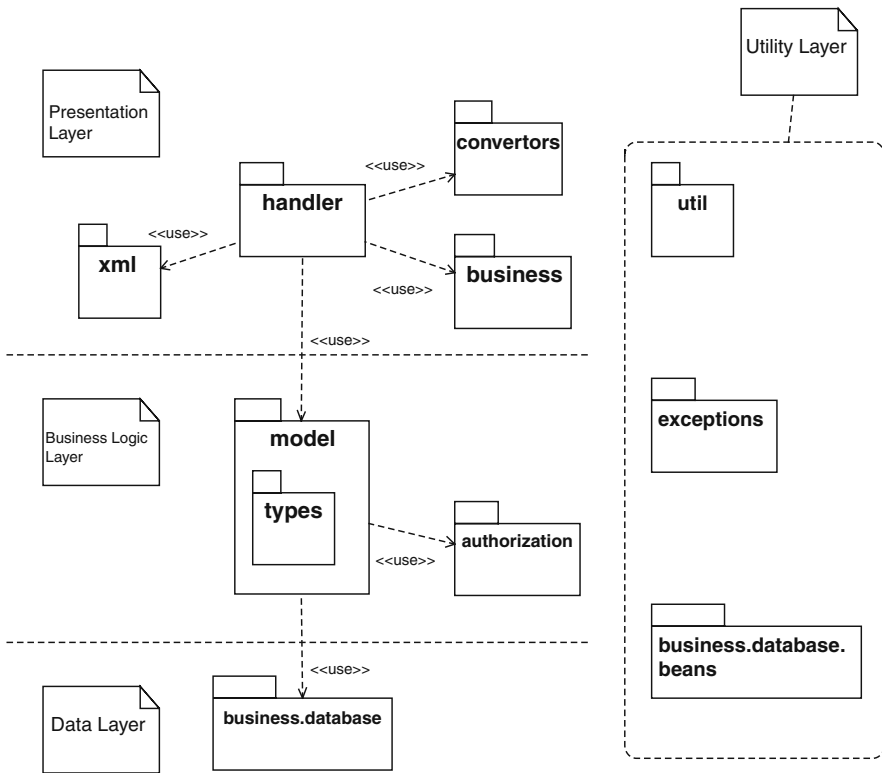


Fig. 19.4 Domain entity services package diagram

This concludes the theoretical side of the introduction into the system’s implementation and leaves us with the practical side. The example in the next section covers the practical side.

Table 19.1 Package descriptions for the domain entity services

Package	Description
convertors	Classes for converting objects from a class convenient for XML representation to a class convenient for business logic handling (<i>CommunityConvertor</i> , <i>CompetenceConvertor</i> ,) (steps 7 and 10).
xml	The data object classes used by Castor for converting XML messages to Java objects and vice versa (steps 6 and 11).
handler	Contains the <i>MainDispatcher</i> and <i>MainFilter</i> control classes and the servlets for receiving and returning XML messages (<i>CompetenceProfileServlet</i> , <i>CompetenceServlet</i> , ...) (steps 2, 3, 4, 5, 12 and 13).
authorization	Checks whether a request should be allowed or rejected, using the <i>Authorizer</i> class (step 8).
model	Contains the main business logic classes, contained in <i>CompetenceProfileManager</i> and <i>CompetenceManager</i> , (step 9)
types	Contains a number of common reference types. Example: <i>GroupType</i> , containing values Admin, Member and User.
business	Helper classes for the Dozer library. Note: the package name is misleading, because it resides in the Presentation Layer.
business.database	Generic classes for persisting data to the database and retrieving data from the database.
util	General utility classes (in version 1.0 there's exactly one class – <i>PropertyLoader</i> – available), which can be used in all layers.
business.database.beans	The data object classes used by Hibernate to persist Java objects to the relational database and to retrieve objects from that database. The package is placed in the Utility Layer because these data classes are also used for transporting/using data in the Business Logic Layer.
exceptions	The custom exception classes for TENCompetence exceptions. These exception classes can be used in all layers.

19.5 Example: Creating a New Service

In this example, we follow the 13 steps to create a new domain entity service. One example is insufficient to grasp all details but it should provide a good impression of the changes required to develop a new service. The example, the guidance given by existing source code and the source code's Javadoc (Sun 2008) should provide a good starting-point for developers to create their own services.

We will create a new service. It will retrieve a person's competences, including their attained proficiency level for each of them. Before starting to code, the structure for the XML message is defined. The combination of a competence and its attained level has already been used in the domain entity services (as element within the competence-level element of the competence profile message). Reusing and extending this structure results in the XML message of Example 19.14.

Authentication (step 1) and opening a transaction (step 2) are performed by use of generic code. Defining a URL for the new service and mapping it to a servlet (step 3) is our next coding action. Based on the REST principles, we take the user URL

```

<?xml version="1.0" encoding="UTF-8" ?>
<attained-competences xmlns="http://www.tencompetence.org/api/v1.0">
<competence-level>
  <level>6</level>
  <competence>
    <link type="competence" href="http://our.server.org/↵
      TENCServer/com1.community/c1.competence">
      Project Management
    </link>
  </competence>
</competence-level>
<competence-level>
  <level>5</level>
  <competence>
    <link type="competence" href="http://our.server.org/↵
      TENCServer/com1.community/c2.competence">
      Research funds raising
    </link>
  </competence>
</competence-level>
...
</attained-competences>

```

Example 19.14 Attained competences

and add a postfix for the attained competences. Many of the existing domain entity services use a community URL as prefix, but this is not appropriate for our new service. We use the user URL as a prefix to express how the retrieved information relates to one specific user and might be spread over multiple communities (which is the case when a user attained competences in more than one community).

Example 19.15 shows the URL's regular expression. An example of an actual URL is `http://our.server.org/user/freud.user/attained-competences`. From Example 19.15 one can see 'AttainedCompetences' is the (new) servlet to which the mapping points. As for all Tomcat servlets, the `web.xml` is used to map the logical 'Attained-Competences' name to the servlet class. Because we offer the service for information retrieval, the HTTP GET method is the only method we will implement in the new servlet.

```

AttainedCompetences= ↵
  /user/([^\\\.]+)\\. (user)/attained-competences(\\Z|\\?.*)

```

Example 19.15 URL mapping

The servlet needs to extract the client identifier for the user whose attained competences need to be retrieved (step 4) and map this to a server identifier (step 5). *MainDispatcher* makes the client identifiers available as attributes of the request. The brackets of "[^\\\.]+" indicate to *MainDispatcher* that it contains an identifier and the brackets in "(user)" tell *MainDispatcher* that the previous identifier is for a user entity. The combination of the extracted identifier and entity type uniquely

identifies an entity in the database. The servlet can get the user id from the request and map it, using a lookup utility, to a server user id.

The service is a GET only service and therefore does not have an XML body in the request. This means unmarshalling from XML (step 6) and converting to a business format (step 7) is not needed. Example 19.16 shows the main code for the new servlet.

```
protected void doGet(final HttpServletRequest request,
                    final HttpServletResponse response)
throws ServletException, IOException {

    // Get the client id of the user for which the attained competences
    // need to be retrieved
    final String externalUserId = (String) request.getAttribute("user");

    // Lookup the server id for the given client id
    final Integer userId = ←
        DBUtil.getInternalId(externalUserId, ObjectType.User);

    // Get the logged on user from t
    final DBUser loggedInUser = ←
        (DBUser) request.getAttribute(HandlerUtil.TENC_USER);

    // Retrieve the user's attained competences
    List<DBCCompetenceLevel> attCompetences = ←
        CompetenceManager.getAttainedCompetences(userId, loggedInUser);

    // Translate from business to XML representation format
    final CompetenceLevels xmlAttCompetences = ←
        CompetenceConvertor.toXmlAttainedCompetences(attCompetences);

    // Marshall to XML and return to the client
    HandlerUtil.setResponse(response, xmlAttCompetences);
}
```

Example 19.16 Servlet code

We proceed with authorisation (step 8) and the actual business logic code to perform the request (step 9). The business logic goes as a new method into the *CompetenceManager* class, which is the handler for all business logic related to competences. For simplicity's sake we will incorporate the easy authorisation check as part of the business logic. Combined with a few lines of business logic code we get the code of Example 19.17 as extension for the *CompetenceManager*. In this business logic, we also see an example of the Dozer bean mapper usage.

In the servlet code shown in Example 19.16, the *List<DBCCompetenceLevel>* result is converted to the new *CompetenceLevels* class (from the *xml* package). This is the conversion from business friendly to XML representation friendly format (step 10). It helps the Castor library to convert to XML, as we can map the *CompetenceLevels* class to the attained-competences element in the result message. Example 19.18 shows the code of this very simple class.

```

public static List<DBCompetenceLevel> getAttainedCompetences (final
    Integer userId, final DBUser loggedInUser)
throws TENCDatabaseException, TENCAuthorizationException {

    // Check authorization: allow only retrieval of one's own
    // attained competences.
    if (!userId.equals(loggedOnUser.getId())) {
        throw new TENCAuthorizationException(
            new ErrorMessage("It's not allowed to ... ", null));
    }

    // Retrieve all attained competences, for this user
    final DBManagerInterface dbMgr = DBManagerFactory.getDBManager();
    Criteria crit = dbMgr.createCriteria(DBUserCompetenceLevel.class);
    crit.add(Restrictions.eq("user", loggedInUser));

    List<DBUserCompetenceLevel> attUserCompetenceLevels =
        dbMgr.getObjects(DBUserCompetenceLevel.class, crit);

    // Create a result without the user member
    List<DBCompetenceLevel> attCompetenceLevels =
        new ArrayList<DBCompetenceLevel>();

    final MapperIF dozerMapper = ApplicationSettings.getDozerMapper();
    dozerMapper.map(attUserCompetenceLevels, attCompetenceLevels);

    return attCompetenceLevels;
}

```

Example 19.17 Authorisation and business logic

```

public class CompetenceLevels {
    private List<CompetenceLevel> list;

    public CompetenceLevels() {
        list = new ArrayList<CompetenceLevel>();
    }

    public List<CompetenceLevel> getList() {
        return list;
    }

    public void setList(List<CompetenceLevel> list) {
        this.list = list;
    }
}

```

Example 19.18 *CompetenceLevels* class

Marshalling to XML (step 11) is the last thing we need to work on, because committing the transaction and returning the XML are done by generic code. To enable Castor to do the marshalling, we define a mapping for the *CompetenceLevels* class. *CompetenceLevels*' only member is of a type (*CompetenceLevel*) already mapped before, because it was used within the competence profile service. This makes mapping for our new attained competences service easier. We only need to map the

CompetenceLevels class itself to the acquired-competences element and its contained List to a collection of competence-level elements. Example 19.19 shows how this was done.

```
<class name="org.tencompetence.tencs.xml.CompetenceLevels">
  <description>Default mapping for CompetenceLevels</description>
  <map-to xml="acquired-competences"/>

  <field name="list" collection="collection"
    type="org.tencompetence.tencs.xml.CompetenceLevel">
    <bind-xml name="competence-level" node="element"/>
  </field>
</class>
```

Example 19.19 Castor mapping

This completes our example of introducing a new service for retrieving a person's attained competences. Going through the 13 steps gives a good impression what is involved when adding new services to the Learning Network infrastructure. This example, together with the source code of the existing domain entity services, should provide a good starting point for developing new services.

19.6 Conclusion

The Learning Network infrastructure provides a set of domain entity services, which are essential for a Learning Network. These services have been informed by the TENCompetence domain model, which is described in Chap. 18. This chapter provides useful and practical information for developers who want to implement applications for a Learning Network using the services.

These services are being deployed on a Learning Network server, which can be setup by anyone interested in hosting a Learning Network. Although these servers are autonomous and can be used independently, it is possible to federate them as desired. These services form the backbone of a range of applications that support the professionals in a Learning Network. These applications can access the domain entity services through their RESTful APIs.

Using an example, we explained how an application should use these APIs. This was achieved by discussing the principles of REST in general and describing how they have been applied to the domain entity services in detail. We addressed topics such as discovery of services, fetching of resources, creation of new resources, URL patterns, updating of resources and the multiple representations of resources. Furthermore, the basic principles of authentication and authorisation and how this reflects on the XML bindings of the resources through their canonical, personal and rights sections, have been explained. This information, in combination with the reference API documentation (Lemmers et al. 2008), should enable developers to start implementing their own Learning Network applications using the Learning Network infrastructure.

There may be situations where the domain entity services are not providing the required functionality for an application. In these situations the services can be extended. We described how the domain entity services are implemented and demonstrated by example how these services can be extended with a new service.

Concluding, the information provided in this chapter should provide a good base for everyone wanting to extend or use the Learning Network infrastructure for developing their own Learning Networks applications. We encourage anyone who is interested in doing so, to join the TENCompetence community and share their efforts with others. Please visit the TENCompetence website (<http://www.tencompetence.org>) for detailed information on how to join the community.

References

- Apache Software Foundation: Apache Tomcat. <http://tomcat.apache.org/> (2008). Cited 8 Sep 2008
- Berners-Lee, T., Fielding, R., Masinter, L.: *Uniform Resource Identifier (URI)*. <http://tools.ietf.org/rfc/rfc3986.txt> (2005). Cited 18 June 2008
- Castor: The Castor Project. <http://www.castor.org/index.html> (2008). Cited July 01 2008
- Dozer: The Dozer Project. <http://dozer.sourceforge.net/> (2008). Cited 1 July 2008
- EQF: European Qualification Framework. http://ec.europa.eu/education/policies/educ/eqf/index_en.html (2008). Cited 27 June 2008
- Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., et al.: Hypertext Transfer Protocol (HTTP/1.1). <http://tools.ietf.org/html/rfc2616> (1999). Cited 18 June 2008
- Fielding, R. T.: Architectural Styles and the Design of Network-based Software Architectures. <http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm> (2000). Cited 12 June 2007
- Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., et al.: HTTP Authentication: Basic and Digest Access Authentication. <http://www.ietf.org/rfc/rfc2617.txt> (1999). Cited 18 June 2008
- Hibernate: Hibernate: Relational Persistence for Java and .Net. <http://www.hibernate.org/> (2008). Cited 1 July 2008
- Koper, E. J. R.: TENCompetence Domain Model. <http://hdl.handle.net/1820/649> (2006). Cited 3 May 2007
- Leach, P., Mealling, M., Salz, R.: A Universally Unique Identifier (UUID) URN Namespace. <http://tools.ietf.org/html/rfc4122> (2005). Cited 1 July 2008
- Lemmers, R., Vogten, H., Martens, H., Beauvoir, P., Dicerto, M.: ID3.6 Architecture Design Document. <http://hdl.handle.net/1820/1243> (2008). Cited 26 July 2008
- MySQL: MySQL Database Management System. <http://www.mysql.com/> (2008). Cited 1 July 2008
- Openfire: Website of ignite realtime. <http://www.igniterealtime.org/projects/openfire/index.jsp> (2008). Cited 27 June 2008
- Sun: Javadoc. <http://java.sun.com/j2se/javadoc/> (2008). Cited 7 July 2008
- TENCompetence consortium: *TENCompetence Project*. <http://www.tencompetence.org> (2007). Cited 1 Feb 2006
- W3C: *SOAP*. <http://www.w3.org/TR/soap/> (2007). Cited 1 July 2008

Architecture have been presented and explained. This section focus on the implementation of working applications running on the Learning Network infrastructure. The presented applications are only some examples of how it is possible to harvest and operationalise the underlying services. Other applications, both web-based and rich client ones have been produced in the project, but due to space limitations cannot be presented here. Even more applications, though, have not yet been created but are feasible due to the open and extensible nature of the service infrastructure.

Each of the examples presented in Chap. 19 represents a specific aspect of the infrastructure. We will describe each application on a functional and technical level in their current state of development. Each tool is equipped with a purpose specific functionality that will support learners and providers of learning in a wide variety of operations, covering planning, authoring, adopting, amending, and using learning opportunities for professional competence development either in a formal or in a personal setting.

All tools use the integrated Learning Network services and interoperate with them in various ways to enable user interactions based on pre-defined usage profiles, such as Personal Development Planning (PDP) or Assessment.

The four tools we present in this chapter not only share a common service architecture and domain model, they also follow a common philosophy and common aims in their implementation. The main driver for this is openness and extensibility. To achieve these goals a modular approach has been chosen and business logic been separated from the user interface. The component nature of the tools allows access to additional web services available not only from the Learning Network servers, but also from third parties (cf. Sect. 21.2).

The applications provide context-aware services, exposing the user to the environment most relevant to their choice of action. Ease of use is one of the generic demands put upon each and every implementation in the Learning Network range of tools.

For some of the tools, the underlying data model needed to be extended, for example to accommodate the new PDP entities which required a quite different approach to data ownership and privacy than more community-based services like competence profiles in the PCM.

Modern ways to convey information and interaction to users about e.g. new learning opportunities or the management of their own learning have been implemented in the form of widgets, i.e. little embedded applications that run within another environment. This approach is a very powerful one when it comes to customising your personal learning landscape, but also when it comes to creating new innovative services or extending existing ones.

The PCM is a software client developed as part of the TENCompetence project (Koper and Specht 2007; TENCompetence 2006), a 4 year IST funded initiative. The main purpose of the software is to help people plan and carry out professional competence development activities in formal, non-formal and informal contexts. Typical professional competence development scenarios support persons in updating their knowledge, skills, attitudes within their jobs, functions or interests; people might want to study for a new function or job to improve their current job level or to develop new interests; people might want to reflect on their current competences to identify which functions and jobs are within their reach and as a result might want to improve their proficiency level for specific competences. All these scenarios will almost certainly result in some learning needs. The PCM helps these persons in identifying suitable learning opportunities or creating new ones if suitable ones are missing. All this is done in the context of a community where users can work individually or in groups. These communities effectively act as artefacts for the Learning Networks. People can turn to a community for finding support for learning problems but equally can contribute to this community as well. The Learning Networks created this way are not governed by any central authority and they can be set up by anyone. The creator of a Learning Network is also the owner of the community and determines policies for access to the Learning Network. In Chap. 19, we already discussed this principle of access control through rights that can be set for some of the entities. The general idea is that the PCM should tend to openness whenever possible in order to stimulate active participation and contributions of all community members. The PCM relies on the principles of self-organisation to regulate this process (Hadeli et al. 2003). Several other tools have been created for the Learning Network infrastructure, but the PCM is the main authoring environment for all the domain entities described in Chap. 19.

The PCM software is available as open source on SourceForge at: <http://sourceforge.net/projects/tencompetence/>.

20.2 PCM Implementation

In Chap. 19 the API of the core services of the Learning Network infrastructure was discussed in detail. As shown in Figure 20.1, the PCM provides a user interface for the manipulation of the core entities through their APIs. Besides these core services, a number of additional, more autonomous services are provided by the server as well, such as the forum, rating and message services. Figure 20.2 shows an UML implementation diagram depicting the main components of the PCM. In the remaining sections we will discuss each of these components in more detail and explain what their relationships are and how they contribute to the overall application.

The REST Conduit is a component that is responsible for communication with the Learning Network services using the REST protocol (Fielding 2000). For this purpose it supports the GET, POST, PUT and DELETE methods of REST. A singleton of the *ConduitManager* class is the pivot point for the component. The conduit

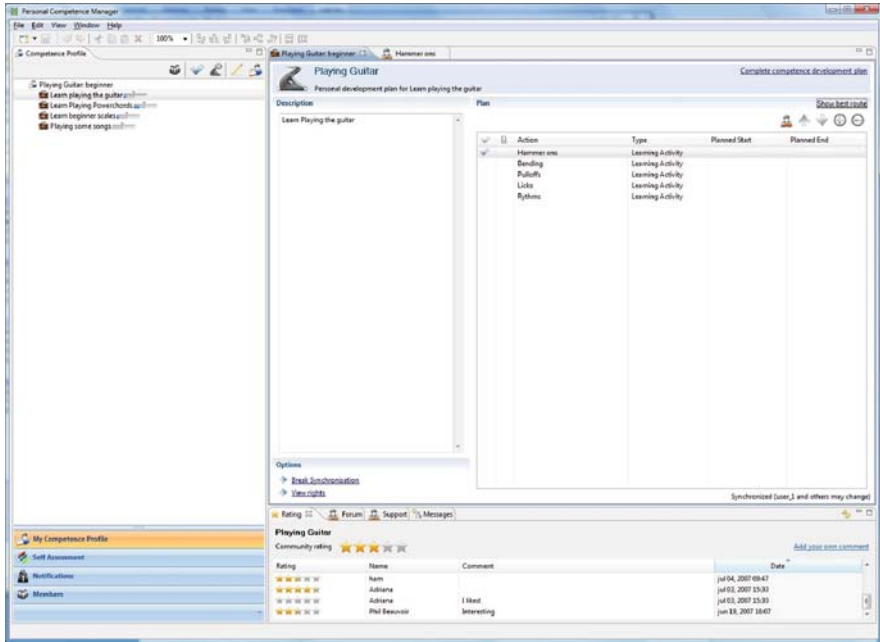


Fig. 20.1 Screenshot of the PCM

manager serves two main functions. First is the factory class for the implementing classes of the *IHTTPProvider* interface. The *IHTTPProvider* interface hides the actual REST implementation and allows different implementations to be used for this purpose. For the PCM we implemented two versions, one based on the Sun HTTP library and the other based on the Apache's HTTP library. The second more important functionality of the *ConduitManager* is taking care of the account management for different REST services.

As explained in Chap. 19, most of the Learning Network services require authentication. Therefore, a user id and password need to be provided when performing one of the HTTP methods. It would be very unpractical if a user had to provide these credentials each time a call was made to any of the services. In fact, we decided for the PCM that a user should only have to provide the credentials once. Therefore the *ConduitManager* is capable of querying and persisting user account information for these services. It does so by maintaining conduit configurations which can be accessed through the *IConduitConfiguration* interface. A conduit configuration associates an URL with a user id and password combination. This association is achieved by matching the protocol, address and port of an URL with the URL of the conduit configurations. The reason for this approach is that authentication is dealt with at a server level. Therefore, the same user id and password should be supplied for all services running on the same server. The protocol, address, and port attributes together uniquely identify an individual server. All conduit configurations are persisted in a single '*conduit.xml*' file stored in the user's home directory. This file

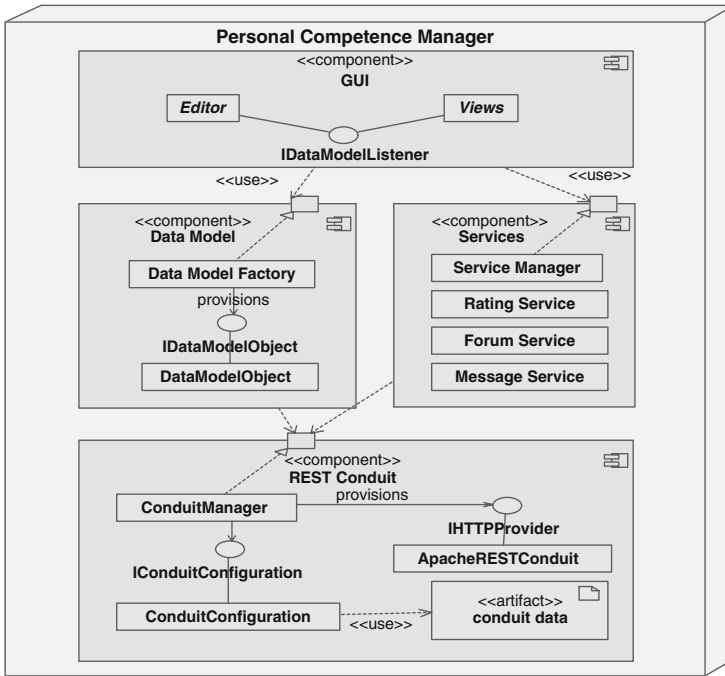


Fig. 20.2 UML implementation diagram showing main components of the PCM

area is protected by the operating system and is only accessible when the user was authenticated by the operating system. Therefore, the PCM does not require a separate authentication process when it is launched.

When a REST method is called for any of the services, the *ConduitManager* will determine whether a conduit configuration exists for this service. If so, this conduit configuration will be passed on to the *IHTTPProvider* together with all other relevant data such as the REST method to user and the URL of the service. However, if no configuration exists for this service, the *ConduitManager* will launch a dialog prompting the user for a user id and password. Only when the service can be successfully accessed with the entered credentials, the *ConduitManager* adds a new conduit configuration for this service. There will be no need to provide any credentials the next time any service hosted by the same server are called. The same prompting mechanism ensures that the PCM even can continue when the ‘*conduit.xml*’ would get lost or is corrupted. The PCM will simply prompt the user to re-enter the credentials for the servers being called. The same applies when the user wants to use the PCM on a different computer and the ‘*conduit.xml*’ was not copied there. Nevertheless, the PCM offers the possibility to export and import the ‘*conduit.xml*’ file.

The *ApacheRestConduit* is performing the actual call to the services when the PCM is installed with its default settings. It uses the conduit configuration data to

authenticate the service call with the server. The data received from the server are simple returned as Strings. If a method fails, an exception is raised. Authentication exceptions are initially intercepted by the *ConduitManager* up to three times, prompting the user to enter new credentials. After that, the authentication exception is re-thrown and error handling is left to the calling party. All other exceptions thrown during the process are also passed on to be dealt with by the calling party.

The functionality described for the Rest Conduit resembles functionality found in web browsers. However there is one difference. The Rest Conduit component also provides a single sign on mechanism for all known servers in the Learning Network infrastructure. See Sect. 19.2 for a discussion about these server federations. Therefore, a web client may want to implement a Rest Conduit component as well, specifically for taking care of this single sign on requirement. Storing the conduit configurations locally can be difficult with web technology due to security constraints. Therefore, a server side solution for this problem is probably a better idea. Of course, this solution could be a proprietary one, but using emerging standards such as OpenID (OpenID Foundation 2008) are probably better alternatives.

The *ConduitManager* acts as a conveyor for the domain entities on the server and the client; it is not aware of the actual data being conveyed. This transformation is the responsibility of the *Data Model* component. This component is capable of serializing Java objects to XML and de-serializing XML into Java objects. A singleton of the *DataModelFactory* class is the pivot for this component. Through this factory class, data model objects implementing the *IDataModel* interface, can be created and fetched. These entities almost corresponds for hundred percent with the domain entities we have discussed in Chap. 19. These entities are not depicted in Fig. 20.2 in order to avoid cluttering up the diagram. Instead these data model objects are represented in the diagram by their common parent class: *DataModelObject*. There are several very good reasons for maintaining a local data model of the domain entities and next we will provide some arguments.

One of these reasons for having a local data model was already mentioned in the previous paragraph: marshalling and de-marshalling of the raw XML data for each of the entities. Any component accessing the domain entities should not be bothered with their XML binding. Instead, the manipulation of these entities is achieved via their corresponding data model objects and methods. Each data model object has a corresponding interface class, exposing the methods for manipulating the object and hiding details about marshalling and de-marshalling.

This marshalling and de-marshalling should also take care of some of the implicit requirements of the REST protocol. For example, when updating an entity via the PUT method, only modified fields should be included in the XML resource being sent to the server. For this purpose all data model objects must keep track of the state of each individual field with respect to the last known state on the server of that field. The application using the data model, in our case the PCM, determines if and when the entities on the server should be updated. The simplest approach would be updating an entity as soon as one of the fields of the corresponding data model object has been changed. However this can have serious performance repercussions. Suppose, we have a dialog that allows all fields of a data model object to be changed.

If we update the server entity for every field this could result in many separate calls, while in principle one call would suffice for this case. Therefore the PCM keeps track of all changes and commits them only when requested to do so. In most cases this update moment is triggered by a user interface component. For example, in the case of dialogs the update occurs after the OK button is pressed; other user interface components such as editors cause the changed entities to be updated after the user have given an explicit save command.

The previous remarks touched upon another reason for having a data model, which is performance. Because the PCM wants to provide a good user experience, the user interface should be responsive. Although fetching entities from a Learning Network server is generally speaking not slow, it can influence the responsiveness of the overall application if these server accesses are too numerous. Therefore, the PCM tries to limit the number of server accesses. The *DataModelFactory* implements two mechanisms to achieve this, which are somewhat related. The first mechanism is a cache. Each time an entity is accessed it is placed on the cache. The URI of these entities are used as a key for retrieving these entities later on. Whenever an entity is requested, the *DataModelFactory* will query the cache and return the corresponding data model object if this entity was found. This of course avoids unnecessary calls to the Learning Network server. Another mechanism implemented by the *DataModelFactory* to increase performance is a form of lazy loading. This lazy loading uses the fact that each entity may have different representation states. Examples of multiple representational states are typically found when using the list services. See Chap. 19 for further details on this REST principle. Although an entity may have multiple representational states, their originating identity can always be established. This is not a characteristic of REST as such, but rather of the *link* element. The *DataModelFactory* is therefore capable of mapping different representational states of one entity to the same data model object. However, not every representational state of an entity contains the full data set for that entity. Therefore, a data model object representing such an entity can be initialised with only partial data. Because each data model object contains the URL of the full representation of that entity, it is possible to refresh the entities content when needed and completing any data fields that may be missing. However, this is only done when an incomplete data field is accessed for the first time. This way unnecessary server access can be prevented because very often these incomplete fields are never accessed in the course of a client session. For example, when retrieving the list of all competence profiles it suffices to only have their partial data. Only after the user has selected a particular competence profile, additional data is required for the selected competence profile.

Another argument for having a local data model for the server entities can be found in the REST protocol itself and was already touched upon. REST makes clear distinctions between creation, updating, fetching and deleting of entities. Especially the distinction between creation and updating can be cumbersome and should be encapsulated from developers using the API. A data model object does this through a specific state attribute. Data model objects can be created in two manners. First they can be retrieved from either the cache or the server, using the appropriate get method

provided by the *DataModelFactory*. This may, or may not, result in a new Java object being created, but will **never** create a new entity on the server. Alternatively, data model objects can be created using the appropriate create method provided by the *DataModelFactory*. This will always result in a new Java object and will result in a new entity on the server once this object is committed to the server.

Our last argument for having a local data model involves the model view controller pattern (Reenskaug 2003) (MVC). This architectural pattern isolates business logic from user interface considerations. Although the MVC pattern has its origins in rich client applications such as the PCM, it is also often seen in web applications, where the view is the actual HTML page, and the controller is the code that gathers dynamic data and generates the content within the HTML. In all cases it uses a model, which is represented by the data model component. The advantages of MVC patterns become particularly clear when various user interface components are using one or more shared data model objects for their model. When such a shared data model object is modified it is necessary to update all user interface components having this object as model. If we would not apply the MVC pattern, all user interface components would have to keep track of all other user interface components in order to notify them of this change. This is very awkward indeed. Instead, by applying the MVC pattern, all user interface components, which are represented by editors and viewer in Fig. 20.2, register themselves with the *DataModelFactory*. This registration implies that these user interface components will be notified of any changes in the data model. For this purpose they must implement the *IDataModelListener* interface. It is the responsibility of the user interface to decide whether a particular change in the data model has any repercussions for the state of the user interface. This way, user interface components do not have to be aware of each other.

The last component of Fig. 20.2 to be discussed is the Services component. This component contains client side services implementation providing access to additional web services available not only from the Learning Network server but also from other servers. An example of such a services implementation includes a forum service, which accesses forum server; likewise the chat service provides access to the Openfire (2008) instant messaging server. Furthermore, there is a rating service allowing users to rate and to comment individual domain entities.

These services provided by the Services component do not have a graphical user interface, but rather act as an intermediate between the views, which provide this interface and the servers. Just like the *Data Model* component this is an implementation of the MVC pattern. Most of these services use the *REST Conduit* component for communicating to their server counterpart, but some are implementing other protocols.

Some of these services require a background process for updating the MVC model to notify the user of changes. For example, the people service uses this background process to update the availability status of the users listed in the people view. And the message service regularly updates the list of messages by fetching the latest messages from the message service on the server. Using a background process for these asynchronous tasks alleviates the user from manually updating the user interface by clicking a refresh button or a similar option.

The actual content presented by these services is depended on the logged on user and the currently active data model object. This is called the data model object context. Depending on this context a service might be required to show different data to the user. For example, the message service shows all messages either sent to the current user or to the community where this active data model object belongs to. Whenever the user selects another community, the service needs to show the messages belonging to this new community instead of the previously selected one. Which data model object is the active context is defined by the application based on the selection of the user. The application uses the *IDataModelObjectContextManager* to notify other components of changes in this active context. The services need to implement the *IDataModelObjectContextListener* interface and register themselves with the context manager in order to be notified of these changes. Whenever a service is notified of such a change, the service determines whether this new context requires an update of its data and if so, it retrieves the new data from the server before presenting it to the user.

References

- Fielding, R. T.: Architectural styles and the design of network-based software architectures. <http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm> (2000). Cited 12 Jun 2007
- Hadeli, P., Zamfirescu, C., van Brussel, S.-G. B., Holvoet, T., Steegmans, E.: Self-organising in multi-agent coordination and control using stigmergy. In: *Workshop on Self-Organising Engineering Applications* (ESOA 2003). http://esoa.unige.ch/esoa03/papers/esoa03_7c.pdf (2003). Cited 19 Jul 2007
- Koper, R., Specht, M.: TenCompetence: lifelong competence development and learning. In: *Competencies in Organisational E-Learning: Concepts and Tools*, ed. by Sicilia M. (Idea Group Inc., Hershey 2007) pp. 230–247
- Lemmers, R., Vogten, H., Martens, H., Beauvoir, P., Dicerto, M.: ID3.6 Architecture Design Document. <http://hdl.handle.net/1820/1243> (2008). Cited 26 Jul 2008
- OpenID Foundation: <http://openid.net/> (2008). Cited 16 Jun 2008
- Openfire: <http://www.igniterealtime.org/projects/openfire/index.jsp> (2008). Cited 27 Jun 2008
- Reenskaug, T.: The Model-View-Controller (MVC), Its Past and Present. In: Java ZONE conference 2003 OSLO. http://heim.ifi.uio.no/~trygver/2003/javazone-jaoo/MVC_pattern.pdf (2003). Cited 16 Jun 2008
- TENCompetence: <http://www.tencompetence.org> (2006). Cited 26 Jun 2008
- Vogten, H., Koper, R., Martens, H., Bruggen, J. v.: Using the personal competence manager as a complementary approach to IMS learning design authoring. In: *Interactive Learning Environments* 16(1), 83–100 (2008)

has also lead to more active and involved society members, who are increasingly more demanding regarding their personal goals and developments. Individuals are regularly confronted with question such as: are my competences still up-to-date?; is my current job still satisfying and challenging enough?; in what directions can I change my career?; can I improve myself?; what other opportunities do I have? These types of question are not necessarily work related, although they often are.

In this chapter we will present an application, the Personal Development Planner (PDP), which supports the professional in reflecting on the aforementioned questions. After a brief overview of the functionality provided by the PDP, we discuss how we used and extended the Learning Network (TENCompetence 2006) infrastructure to implement this application.

The PDP supports the professional by monitoring his progress in a selected competence profile, which fits best with their goals. The PDP not only places the professional on this competence profile but also assists in acquiring any missing proficiencies for this profile. This is accomplished by letting the professional select specific learning activities, which help them in acquiring these missing proficiencies.

When the PDP is launched, the user must define a goal. This is achieved by selecting a competence profile that fits best with their goals. These competence profiles are defined in the context of communities as was discussed in Chaps. 19 and 20. The motivation for selecting this competence profile can be further refined by selecting one of the predefined values ('achieve a position', 'keep up-to-date', 'upgrade') or by entering a free text value.

The PDP also offers the possibility to enter a detailed description for the goal. It is possible to blog progress and experiences gained in the process of achieving the goal. These blog entries are strictly private unless they are explicitly shared with other users who are developing themselves in the context of the same competence profile. In principle all data entered in the PDP are strictly personal and not accessible by other users, unless access is explicitly granted by the user. A user may have defined multiple goals and he can work on them simultaneously by opening them in separate tabs.

After the goal has been specified, the user must be placed on the selected competence profile. Figure 21.1 shows a screenshot of the PDP tool with the self-assessment tab activated.

A self-assessment for a competence profile involves an estimation of the acquired proficiency level for each individual competence of this competence profile. Figure 21.2 shows all competences, including their folder structure, on the panel on the left hand side of the screen. Indicators show the estimated levels as green bars, and the remaining required level as grey bars. On the centre screen, the details of the selected competence are presented. This includes a descriptive text for the competence; a slider that allows setting the estimated acquired proficiency level for the competence; and a collection of evidences supporting the claim for having acquired the competence at the required level. The level indicator varies from zero through to eight representing the proficiency levels as they have been defined in the European Qualification Framework (EQF 2008). A self-assessment test may have been asso-

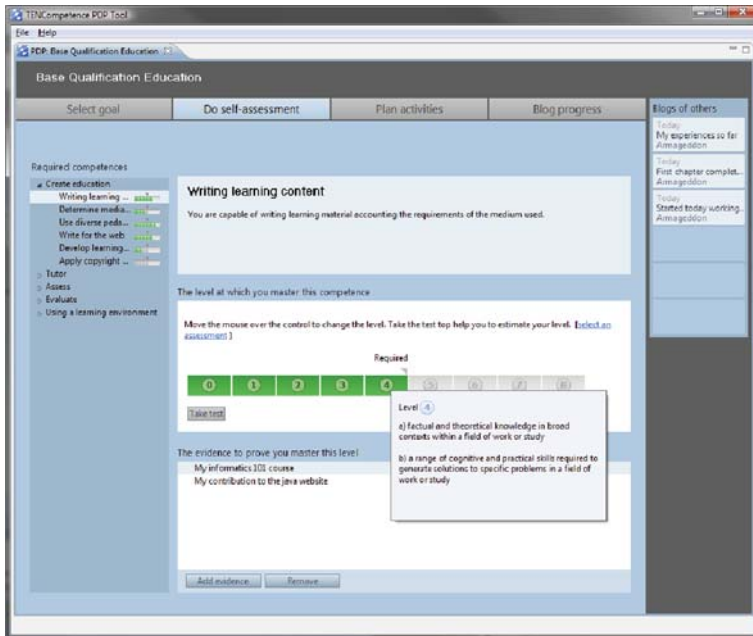


Fig. 21.1 Screenshot of the PDP with the self-assessment tab active

ciated with a competence, which supports the user in determining the acquired level for this competence. The creator of the competence profile can assign, or reassign, these self-assessment tests to the individual competences of the profile. Such a test involves scoring all inquiries into the past experiences of the user on a Likert scale. Based on this score, the PDP will set the assessed proficiency level for the competence. Each competence has a target proficiency level for the competence profile; this implies that the user must attain the competence at this level or at a higher level to meet the requirement of the competence profile. A small triangle shaped marker indicates this threshold. If the outcome of the self-assessment suggests that the competence is mastered at the required proficiency level, or higher, the user can add evidence supporting this assessment. Each evidence entry consists of free text and optional attachments substantiating the claim.

After the self-assessment it is most likely that one or more competences are not attained at the required levels. The user can now move to the next tab that helps with the choosing and planning of the learning activities that will help acquiring the required proficiency levels for these competences. The user may now generate a personal plan for this competence profile by pressing the generate plan button, which automatically adds learning activities to the plan. These activities are selected from the competence development plans available for each competence. If desired, the user may modify these learning activities by removing or adding learning activities. It is also possible to create complete new learning activities and add them to the list. When new learning activities are added to the list, the user must select to which

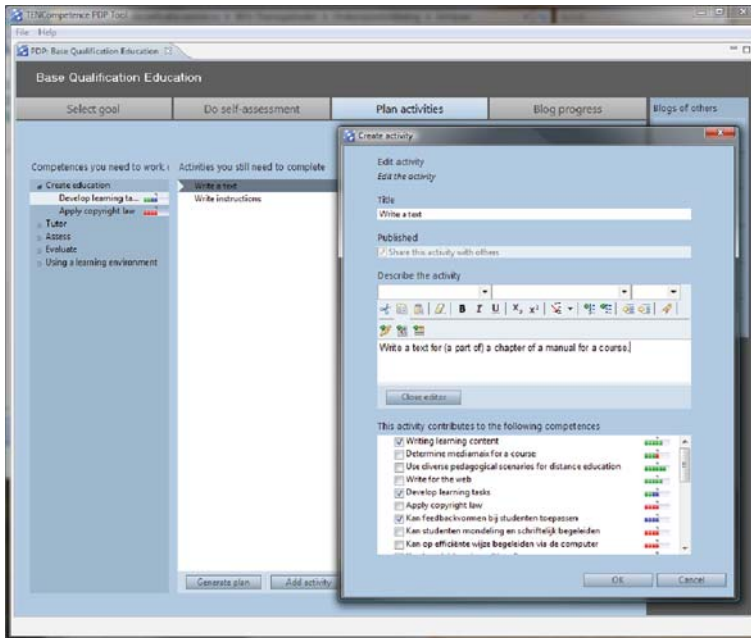


Fig. 21.2 Screenshot of PDP: adding a learning activity in the plan activities tab

competences they contribute. Figure 21.2 shows a dialog for creating a new learning activity.

The left-hand side panel of the screen is showing the competences that have not been acquired so far, and therefore require the planning of learning activities.

Therefore this list is a subset of the complete list of competences for the competence profile. Furthermore, this indicator has two states for showing the gap with the required level: blue bars indicate that learning activities are planned that are targeted towards this competence; red bars indicate that no learning activities for this competence are part of this plan. Once the user has planned all learning activities he can proceed to the next tab dealing with reporting of the progress on the learning activities.

Figure 21.3 shows the progress tab of the PDP. The planned learning activities are shown on the left hand side panel of the screen. The centre panel provides details about the selected learning activity. This includes a descriptive text of the learning activity, which may include hyperlinks to external learning environments such as Moodle (2006) and CopperCore (Vogten and Martens 2008) and a possibility to blog progress with respect to this learning activity. It is possible to mark a learning activity as being complete, which causes it to disappear from the list of learning activities on the left hand side panel. However, through a history mode toggle button, it is always possible to review the full set of learning activities started in the context of this competence profile.

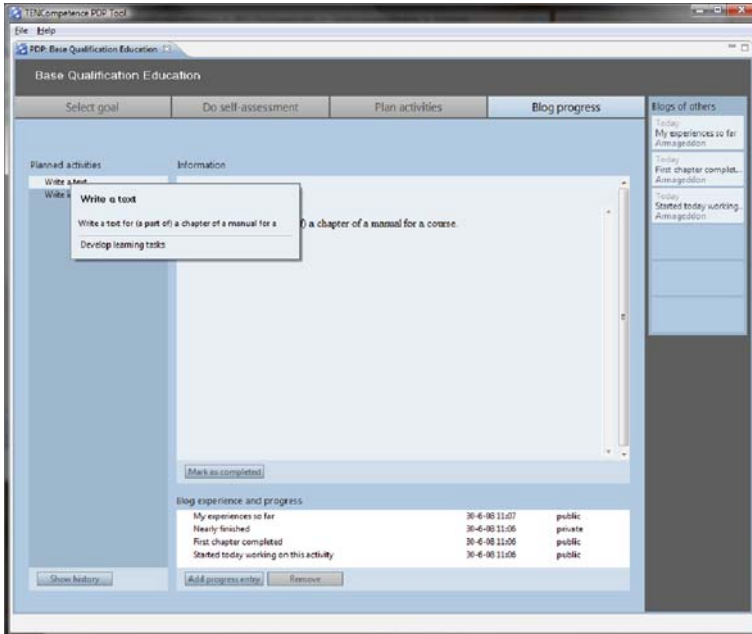


Fig. 21.3 Screenshot of PDP with the blog progress tab active

Finally, it is possible to keep track of progress within the learning activity through blogging. These blog entries can be kept private, or shared with other community members studying the same learning activity. Blogs entries of others are syndicated through feeds and are shown in a separate widget on the right hand side of the screen; they can be consulted any time and could provide valuable insights to other users. After completing all learning activities, the user may place himself again on the competence profile by re-doing the self-assessment.

The process of selecting the goal, taking the self-assessment, planning the appropriate learning activities, and blogging progress have been presented as sequential steps, which they initially certainly are. However, during the process the user may decide to use the options more dynamically going up and forth between the different tabs. For example, reading the blog of others could suggest to the user that it is better to select a different learning activity for some competence.

The PDP tool can be used in different contexts, both on or off the job, formal, informal and non formal. Therefore the PDP does not make any assumptions about the context and gives the user full control of his own data and lets him decide whether to share data captured by the PDP with others. In the job context, the data could be shared with a human resource manager in order to plan and track professionalisation on the job. In an informal situation, the PDP can be used as a pure

personal tool to reflect on progress towards these goals. In this case data will be kept private.

21.2 Some Implementation Details of the PDP

The PDP is implemented as an eclipse rich client application using the Learning Network service infrastructure (Lemmers et al. 2008) as was discussed in Chap. 19. The technological approach for the PDP is very similar to that of the Personal Competence Manager (PCM), which was described in Chap. 20. In fact, the PDP and the PCM share some of the same base components such as the conduit framework and data model. The PDP reuses domain entities such as competence, competence profile, competence development plan, learning activity, assessment activity. Moreover, the PDP depends for the creation of some of these entities on the PCM. However, these domain entity services provided by the Learning Network infrastructure had to be extended with some new entities for the PDP and some existing entities were enhanced. Additionally a blogging service was added to the infrastructure in order to give the possibility to blog progress in a particular context. The remainder of this chapter will focus on these services and we will describe how they were added to the Learning Network infrastructure.

Figure 21.4 shows the UML class diagram of the entities that have specifically been added to the infrastructure for the PDP.

The Personal Development Plan is the core entity of these new PDP entities; every user has zero or more of these plans where each plan is representing a specific goal of the user, which is operationalised through a competence profile; a plan references a list of learning activities, which have been selected to attain all compe-

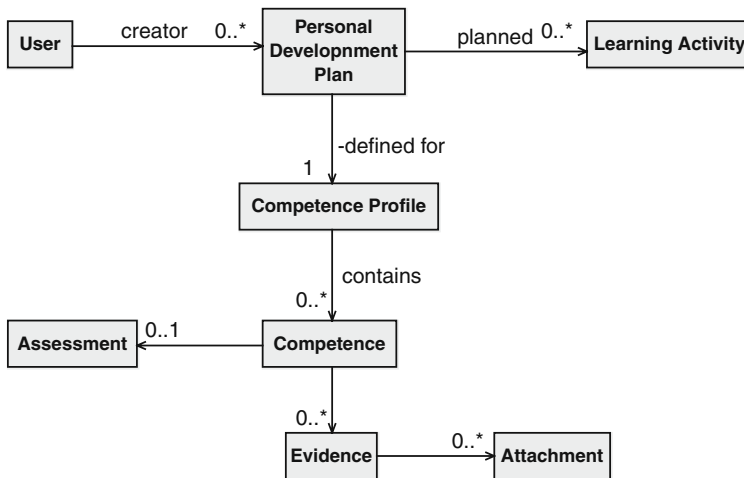


Fig. 21.4 UML class diagram of PDP extensions to the standard entity services

tences of the competence profile. A competence can optionally have an assessment supporting the user making the self-assessment for the competence. Finally, each competence can reference one or more pieces of evidence, which in turn can reference zero or more attachments.

These new entities and their corresponding services have been implemented according to the guidelines described in Sect. 19.4; all services are implemented using a RESTful interface; the XML binding for the resources are divided in canonical, personal and right sections; all relations are implemented through references which use the *link* element for referencing the full representation of other entity elements. As a result, the PDP services could simply be added to the existing domain services and will be part of the next release of the Learning Network infrastructure. Because the new PDP services are also using REST, the standard HTTP conduit described in Sect. 21.2 could be used without any modifications. The data model was extended with Java representations of the new PDP entities. Their implementation was very much in line with the implementation of the existing data model objects for the domain entities.

However the PDP entities are different to the domain entities in some aspects. First of all, the treatment of the default rights for PDP entities differs: access to the PDP entities is restricted to the owner of these entities, while the domain entities are shared by all community members by default. This automatically brings us to another major difference. The PDP entities are not directly defined in the context of a community. Instead they are much more closely related to the owner and the relation with the community is rather defined indirectly via the competence development plan.

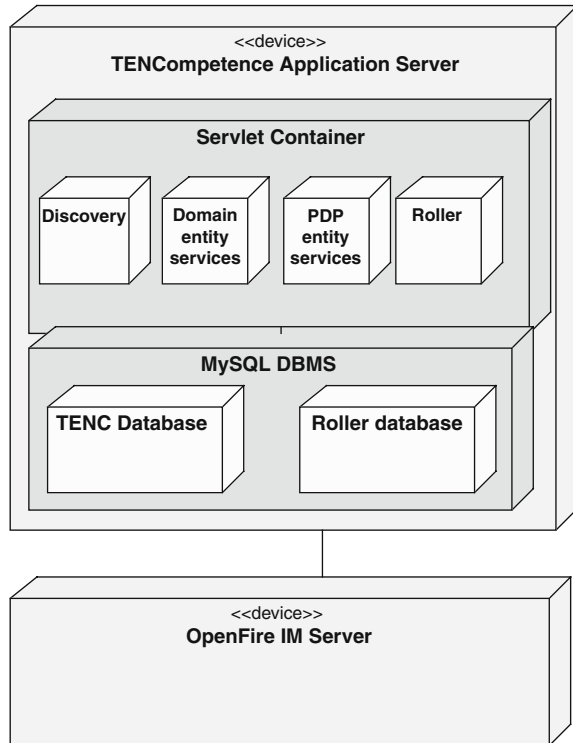
The other service incorporated specifically for the PDP is the Roller blogging server (Johnson 2008). Roller is a Java implementation for a multi user blogger; it allows the creation of personal and group blogs and Roller allows blogs to be aggregated through the site aggregator feature.

The Learning Network infrastructure had to be extended to support the Roller server. Roller is implemented through Java servlets and can therefore be deployed on the Tomcat application server, which is part of the Learning Network infrastructure. Figure 21.5 shows a modified UML implementation diagram of the Learning Network infrastructure that includes the changes required for supporting the PDP services. The diagram shows that Roller is deployed on Learning Network application server, side by side with the existing services and the new PDP services. A Roller database is added to the MySQL DBMS.

However, integration does not stop with deploying the Roller blogging server. The Learning Network server synchronises user account management with Roller and ensures the creation of the personal blogs for every user of the PDP. For this purpose Roller provides a special API, the so called Roller Admin API; it is a RESTful API providing programmatic access to commonly used administrative tasks such as user account management. This API is very similar to another API supported by Roller: the ATOM publishing protocol (Gegorio and de hOra 2008).

The ATOM publishing protocol is a simple HTTP-based protocol for creating and updating web resources. The PDP uses this protocol to publish and update the

Fig. 21.5 UML implementation diagram showing Roller integration



blog entries to the Roller server. Because the ATOM publishing protocol is widely supported, it is possible to replace Roller with another blogging server without having to make changes to the PDP application. However, this is not the case for the Learning Network server, which uses the proprietary Roller Admin API. A common standard is momentarily lacking and proprietary solutions are for now the only alternative.

The PDP automatically tags the blog entries by adding the context of this entry. This context can be the competence profile in general, or a specific learning activity. These tags allow the PDP to fetch specific blog entries from a user, but also from all users working either on the same competence profile or learning activity. This way it is not only possible to show the relevant blog entries from the user, but also from the community. Fetching the blog entries is achieved via syndication of web feeds.

Roller offers this syndication through the Atom Syndication Format, which is an XML language used for web feeds (Nottingham and Sayre 2008). The PDP reads the web feed of an aggregated blog, which combines all blog entries of all Roller users. This feed is filtered on the correct context tags, ensuring only appropriate entries are presented to the user.

References

- EQF: European Qualification Framework. http://ec.europa.eu/education/policies/educ/eqf/index_en.html (2008). Cited 27 Jun 2008
- Gegorio, J., de hOra, B.: The Atom Publishing Protocol. In: IETF. <http://tools.ietf.org/html/rfc5023> (2008). Cited 27 Jun 2008
- Johnson, D: Apache Roller. <http://rollerweblogger.org/project/> (2008). Cited 16 Jun 2008
- Lemmers, R., Vogten, H., Martens, H., Beauvoir, P., Dicerto, M.: ID3.6 Architecture Design Document. <http://hdl.handle.net/1820/1243> (2008). Cited 26 Jul 2008
- Moodle: <http://moodle.org/> (2006). Cited 26 Jun 2008
- Nottingham, M., Sayre, R: The Atom Syndication Form. In: IETF. <http://tools.ietf.org/html/rfc4287> (2008). Cited 27 Jun 2008
- TENCompetence: <http://www.tencompetence.org> (2006). Cited 26 Jun 2008
- Vogten, H., Koper, R., Martens, H., Bruggen, J. v.: Using the personal competence manager as a complementary approach to IMS learning design authoring. In: *Interactive Learning Environments* 16(1), 83–100 (2008)

ReCourse was developed as a successor to the RELOAD Learning Design Editor (RELOAD LD), in order to provide a more accessible way of editing Units of Learning (UOLs). Following discussions with users of RELOAD and observations of workshop sessions the design team identified three principal challenges which were faced by users. Firstly, the structure of the specification is complex, and a tree interface does not support users in holding the whole structure in their minds. Moreover, unlike simpler IMS specifications, the values given to elements may be interdependent, and may affect the behaviour of other elements in the UOL when it is run. This complexity is not well suited to the tree metaphor which lay behind the RELOAD interface. Secondly, the services which authors could include in their Units of Learning were very limited, were often not available in the runtime environment, and were not represented in the authoring tool. Thirdly, the editor was not integrated into a workflow. As a result it was not easy to preview, publish, set up and run a UOL.

The goal adopted by ReCourse was to address these three barriers, and so to create an application which could be used by anyone with a conceptual grasp of pedagogical modelling, and basic computer skills (e.g. the ability to construct a simple website).

22.2 The Development Process

Building on the experience of developing RELOAD LD the design team defined three main design goals for the new LD Authoring tool. The first of these was to simplify understanding of the specification, by clarifying its structure and hiding the XML file format. Secondly, the team sought to provide an authoring interface appropriate for users who are not technical experts. The initial feature list in this respect included user-friendly and configurable language and terms; activity flow diagrams and editing; re-use of UOLs and UOL components; and more intuitive presentation of *role parts*. Lastly, it was decided to restructure the application. The new structure would provide a modular plug-in architecture with which to support an integrated author workflow, and the possibility of supporting different levels of complexity of editing, e.g. through templates and wizards.

From this starting point a participatory design process was established, involving four users of RELOAD from Liverpool Hope University, with less intensive contributions from other users. This was not simply a matter of collecting user needs, as a dialogue was maintained with the user group throughout the development process.

22.3 The Functionality of ReCourse

ReCourse is a generic LD authoring tool that at the time of writing supports IMS Learning Design Level A in both data model and user interface, and LD Levels B and C internally in the data model. An interface for editing Levels B and C is

planned for the next phase of development. It is fully compliant with IMS LD, and aspires to being a reference implementation of the specification. ReCourse creates *Units of Learning* (UOL) an IMS Content Package that contains an *imsmanifest.xml* file with a “learning-design” element as its organisation and usually contains other assets such as HTML resource files, and sometimes other items such as QTI files. The current version of ReCourse is not constrained to a particular use case or use cases, and this leads to an open-ended workflow in the interface. As a result the tool is more effective if the user has a basic conceptual understanding of IMS LD. The present version, however, is the core of the tool, and its modular architecture means that use-case driven views and editors can be implemented at a later stage.

22.3.1 User Interface

The interface of ReCourse is shown in Fig. 22.1 below.

The Organiser pane to the right acts as a library and supports the management of UOLs, resources and elements of the current UOL.

The UOL is constructed in the Learning Design Editor Window. In ReCourse version 1.5 this is sub-divided into three tabs – Main, Environments and Resources.

In the Main tab to the top left the author builds a flow of activities (i.e. the *method*), by dragging activities onto a grid, and indicating which roles will carry them out. This approach enables the IMS LD concept of *role-part* to be eliminated from the interface, as it can be inferred from the author’s actions in the activities grid. Below the list of *Plays* and *Acts* are links to edit the overview, prerequisites, learning objectives, completion rules and the final packaging of the UOL .

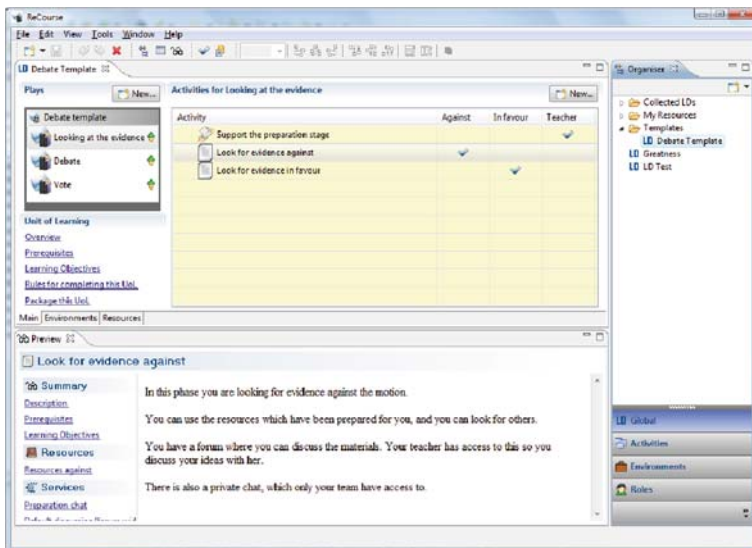


Fig. 22.1 Screenshot of ReCourse

In the Environments tab the author creates new collections of resources and services, drags on components, associates them with resources and edits their properties. These can be associated with *activities*, or can be dragged directly onto the Activities grid in the Main tab.

New files and resources can be created in the Resources tab, and default files and resources created in the authoring process can managed.

Most objects in ReCourse have *properties* which the author sets to affect the content or actions in the UOL, and these are set in the Properties view at the bottom left. For example, the *properties* determine which *resources* are associated with *learning objects*, or which *environments* are used in an *activity*.

22.4 Elimination of Import-Export

The RELOAD LD Editor saves the UOL to local disk storage in an intermediate file format in order that application specific information (e.g. graphical layout) can be persisted alongside the mandatory IMS LD information. This required that UOLs should be imported before they could be edited, and exported before they could be run. ReCourse, however, saves the UOL at all times in the IMS LD XML format, eliminating the whole import-export process. Additional information, such as the authoring tool's user interface notation (colours, positions of graphical elements, descriptions, notes and narratives, etc.) are added directly to the *imsmanifest.xml* file within a name-spaced extension to the "metadata" element. This additional information can then be used by another author using a different instance of ReCourse. The information can be removed from the final manifest file if required, as some systems do not currently support extensions in this way (e.g. CopperCore).

Since the XML file is continuously saved in its final format the *imsmanifest.xml* file is always in a state "ready to go" whether to be added to a zipped Content Package or simply e-mailed to a colleague, providing instant portability. This approach also opens up other possibilities. For example, it would be possible to edit a "live" instance of the XML – perhaps in a raw XML editor module. Similarly it could be leveraged "to view the design as it may finally appear" (Wilson 2005) by applying a stylesheet to the *imsmanifest.xml* file instantaneously. Because the information is a name-spaced extension to the specification, other future tools could be designed which would also read this information.

22.5 Modular Plug-In Architecture

The plug-in architecture for ReCourse is inherited directly from the Open Source Eclipse framework on which it is built, and consequently the core components of the application are modular. This facilitates maintenance, and the inclusion of additional features without impacting on the core code-base. The application is a Java Rich Client Platform (RCP) desktop application which provides cross platform support

with a native user interface. The developers of ReCourse are keen to create not only a desktop client application but also an authoring framework and architecture. The following sections briefly describe the additional plug-ins that have been developed as of the time of writing.

22.5.1 Supporting Workflow with Additional Plug-Ins

A problem identified with RELOAD was that it was too isolated from the author's workflow. The plug-in architecture of ReCourse has been leveraged to address this by including functionality which enables authors to carry out many necessary tasks from within the application.

22.5.2 OpenDocument.net

Two usage profiles are foreseen for a repository in relation to UOLs. Firstly, it can support collaboration in the creation of UOLs, providing somewhere to store, browse and use example UOLs, resources, components of UOLs, part finished or hacked UOLs, etc. Secondly, it provides a place where the UOLs to be used in TENCompetence Competence Development Plans can be stored and referenced. ReCourse currently addresses the first of these with a plug-in that connects to the OpenDocument.net repository. A similar mechanism will be implemented to enable authors to upload to the TENCompetence Fedora repository to cover the second usage profile.

22.5.3 Publishing and Provisioning

Units of Learning can be uploaded and published to a CopperCore server from within the ReCourse editor. UOLs on the server can then be managed allowing "runs" to be created and populated with real users. All entities may be created and deleted. The UOL is firstly saved to a zip file as part of the Package & Publish wizard or it may already exist as a zip file in Content Package format. Given the URL of the CopperCore server the zipped UOL is uploaded to that server.

22.5.4 QTI Editor

A simple editor has been developed which creates questionnaires and tests compliant with a subset of QTI 2.1 (as described in Chap. 23). This optional component was not developed by the core developers, and so demonstrates that the plug-in architecture is effective. When the interface for Level B is available in ReCourse properties defined in the QTI editor will be aligned with those defined in the UOL.

22.5.5 Authoring Services

IMS Learning Design describes and supports four types of abstract services – send mail, conference, index search and monitor. This set is rather restricted, and the integration of services at runtime had not been satisfactorily resolved. In order to remedy this TENCompetence has extended the CopperCore runtime system with a widget server (called Wookie), which provides an extensible set of services at runtime (for example chat, forum, and vote). In order for the designer to be aware of what widget services are available for a given server, and to provide a convenient method of adding these services to an Environment in ReCourse, a plug-in provides a link between the editor and the widget runtime system. The user provides the URL of the Wookie server and ReCourse then shows the available services on the environments palette, making them available for inclusion in the UOL. This innovation does not compromise the IMS LD compliance of the UOL, as the only addition is that the *parameters* attribute of a service is allocated a string such as *widget=chat*. This indicates to the server at runtime that the given service requires an instantiation of a particular widget. SLeD is distributed as part of the integrated TENCompetence runtime environment, together with default widgets for frequently used services such as chat, forum and vote. A convention has been established to refer to these services (e.g. a synchronous conference should be referred to as a *chat*, rather than a messenger). The Wookie administrator can include a number of different *chat* widgets, and designate one as a default. If a specific *chat* widget is requested by an author, then the server will provide this if it is available. If not, the default chat widget will be provided. This arrangement maintains interoperability of runtime environments and also supports flexible services.

22.6 Progress Towards an Infrastructure for IMS LD

Wilson has proposed an “Architecture to Support Authoring and Content Management with Learning Design” (Wilson 2005) which remains the key analysis of this topic. We now identify the progress which has been made towards the infrastructure which he proposed, and discuss some of the outstanding aspects. Wilson (*ibid.* p41) identifies the main tasks foreseen for an IMS LD authoring system. We now summarise progress towards these desired functionalities in Table 22.1.

Wilson also notes that “it may be possible to create a framework in which the various components of the editor can be “plugged in” as one approach to collaboratively developing an editor” (*ibid.* p. 51). This proposal is fulfilled in the modular, plug-in approach adopted by ReCourse, as demonstrated by the inclusion of the editor for assessments.

The Valkenburg reference architecture for Learning Design (Wilson 2005) also specifies that the LD editor should integrate with other packages in the architecture. These packages amount to the following levels of access:

Table 22.1 Progress towards the functionality identified by Wilson

Identified by Wilson		How addressed in ReCourse
1	Constrain the variety of learning designs.	Some output validation provided by the checker module. Modification of the behaviour of the editor is not supported.
2	Create, edit and store learning design templates.	The ability to create and edit a new UOL based on an existing UOL provides basic template functionality. The template can be stored and accessed using the OpenDocument.net plug-in.
3	Create and edit learning designs.	This is the principal functionality of ReCourse
4	Edit presentation of learning designs.	The use of stylesheets is not supported. The capability to select preferred and default widgets is a step in this direction.
5	Discover and add materials to learning designs.	Materials can be discovered using OpenDocument.net, and through the TENCompetence LearnWeb2.0 repository.
6	Aggregate learning designs.	Not implemented in ReCourse
7	Create, edit and store materials.	ReCourse includes a built in rich text editor, an organiser for bookmarks, and integration with the OpenDocument.net repository
8	Test learning designs.	The validity of learning designs can be tested using the checker.
9	Store learning designs in a repository.	Integration with the OpenDocument.net repository
10	Search and retrieve a learning design from a repository.	Integration with the OpenDocument.net repository

1. Access to repositories for search, storage and retrieval of UOLs and materials.
2. Access to the runtime elements of the overall architecture.
3. Access to the provisioning elements of the overall architecture.

As described above, integration for points (1) and (2) has been substantially extended in ReCourse, with the author having access to repositories, runtime set up and information about runtime services from within the application. Point (3) is currently being addressed within the wider context of TENCompetence, where processes and systems are under development to link competence development plans with the TENCompetence repository and the Learning Design runtime system, although it is not yet clear what functionality will need to be added to the authoring component to support this (if any).

As can be seen from the above analysis, ReCourse fulfils most of the requirements of Wilson's architecture. We now review those that remain outstanding (using the numbering in the above table)

- (1) *Constrain the variety of learning designs.* Authors who wish to constrain the behaviour of the editor itself are directed to the CRT RELOAD tool developed in the TELCERT project (O'Neill et al. 2005). The TELCERT tools also pro-

vide more sophisticated validation and conformance testing than is available in ReCourse, but do not offer the same degree of usability.

- (4) *Edit presentation of learning designs*. The use of stylesheets is not supported in run-time systems at present, and so there is little to be gained by introducing them in the authoring tool. Nevertheless this remains an interesting proposal.
- (6) *Aggregate learning designs*. Another UOL can be referenced by its URI within the LD, and indeed it is possible for a *role part* to consist of a *role* doing another UOL rather than an *activity*. It is desirable that ReCourse should have an interface to achieve this.

Wilson also proposes two interfaces which support two sets of functions: (a) “creating pedagogic scenarios, defining the flow of activity along with the various branching conditions. . .” and (b) “populate a design with specific resources and services” (Wilson 2005, p. 50).

At the time of writing ReCourse provides an editor for case (a) (pending completion of the Level B authoring interface). Case (b) is not addressed in the current version, but the Eclipse development environment used provides facilities for adding wizards, Eclipse cheat sheets or perspectives which are well suited to this functionality.

Similarly resources in the Global Organiser are at present ‘bookmarks’ with very basic meta-data support for *type*. This could be extended in line with Wilson’s proposal for a Meta-data Editor used to “tag materials with metadata to facilitate discovery of materials for use within LDs” (ibid. p.49). For example, possible types could include views on repositories for collected full and part UOLs and Templates. Similarly the component parts of LDs could include *activities*, *environments* and entities of various granularity (e.g. *activity structure + environments + level B properties* could constitute such a re-usable entity). Collections of materials a stylesheets could also be made available.

ReCourse 1.5 is by no means the final word in designing usable interfaces for editing IMS LD. Firstly, as can be seen from the above analysis, some desirable functionality remains to be addressed. Secondly a wide range or other alternative interfaces could be created, including the reduction of complexity by constraining the options available to the author. Within this context, however, ongoing work with users indicates that substantial steps have been taken towards an editor for IMS LD which can be used by those who are not technical expert.

References

- Griffiths, D., Blat, J., et al.: Learning design tools. In: *Learning Design: A Handbook on Modelling and Delivering Networked Education and Training*, ed. by Koper R., Tattersall C. (Springer, Berlin, Heidelberg 2005) pp. 109–135
- Griffiths, D., Liber, O.: Opportunities, achievements, and prospects for use of IMS LD. In: *Handbook of Research on Learning Design and Learning Objects: Issues, Applications, and Technologies*, ed. by Lockyer, L., Bennett, S., Agostinho, S., Harper, B. (Hershey, New York 2008) pp. 87–112

IMS Global Learning Consortium Inc: Learning Design Specification Version 1. <http://www.imslobal.org/learningdesign/> (2003). Cited 21 Jun 2008

Wilson, S.: Architectures to support authoring and content management with learning design. In: Learning Design: A Handbook on Modelling and Delivering Networked Education and Training, ed. by Koper R., Tattersall C. (Springer, Berlin, Heidelberg 2005) pp. 41–62

approaches as units of learning. However, LD can not explicitly support various types of assessment. Assessment components within the Educational Modelling Language (Koper 2001), the basis of the LD specification, were excluded when LD was adopted by IMS, because of the existence of IMS Question and Test Interoperability (QTI 2001). QTI is also an open e-learning technical standard, which describes a data model for the representation of assessment item/test and their results. It defines a set of interaction types which can be used to specify basic question types (e.g., multiple-choice, fill_in_the_blank, and slider) and complicated question types through combination. In particular, since the release of QTI v2 (QTI 2006), the integration of QTI and LD is possible at the specification level.

This section describes our work on supporting QTI authoring and the integration with ReCourse. Such an integrated authoring environment makes it possible to specify an integrated learning, teaching, and assessment process at technical level.

23.2 The Benefits to Use QTI

As mentioned above, QTI is an open e-learning technical standard about assessment. In QTI, a question is called an assessment item, which not only describes the question itself, but also contains information about how to handle the responses. An exam is called an assessment test, which represents a set of structured items and the rules to calculate the final score. In addition, QTI introduces an intermediate structuring construction called section, which is used to group individual items or other simpler sections together. QTI also enables the declarative description of many relevant dynamic features of the assessment process. For each structural element it is possible to attach a set of rules used to process learners' responses. These rules can, for instance, trigger the presentation of some feedback and change the order of presentation of the items. Their formulation can be based on a set of control switches, which are also included in the description of the corresponding structural element or their neighbours.

23.2.1 Powerful Expressiveness

QTI v2 is an extensive and sophisticated specification that allows the creation of diverse and complex questions. Almost all classical question types can be supported by QTI. Moreover, it provides sufficient flexibility to grow into the advanced constructed-response items and interactive tasks we envisage as the future of assessment elaborates the assessment items in detail (Almond et al. 2001).

23.2.2 Supporting Interoperability and Reusability

When trying to reuse educational content in different Learning Networks, we can agree on a common exchange format. QTI is an open format developed by a neutral organisation IMS, and not by a single vendor. Therefore, assessment ReCourses and tools compatible to QTI can be shared and reused in Learning Networks.

23.2.3 Integration with Learning/Teaching Processes

QTI v2 provides the possibility to integrate QTI items with a LD unit of learning (UOL). The primary motivation for integrating LD and QTI stems from use cases involving formative assessment and summative assessment using items with traditional question types (QTI 2006). Concretely speaking, an outcome variable of a QTI item can be coupled with a property defined in a UOL, which can be viewed by the participants of the UOL and can be used to control the learning processes.

Note that a QTI item/test is represented in the form of XML. Figure 23.1 shows an inline-choice question with one correct answer. This XML example does not yet contain response processing and showing feedback. That would have made the listing more complicated and longer.

```
<?xml version="1.0" encoding="UTF-8" ?>
<!-- This file was created with the TENCompetence QTI Assessment Item Editor on Sun May 18 23:27:36 CEST 2008 -->
- <assessmentItem xmlns="http://www.imsglobal.org/xsd/imsqti_v2p0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.imsglobal.org/xsd/imsqti_v2p0_imsqti_v2p0.xsd" identifier="inlineChoice"
  title="Richard III (Take 2)" adaptive="false" timeDependent="false">
  <outcomeDeclaration identifier="SCORE" cardinality="single" baseType="integer" />
- <responseDeclaration identifier="RD-d81acfa0-3f41-4da9-91dc-de9c7b2d6920-46" cardinality="single"
  baseType="Identifier">
  - <correctResponse>
    <value>IC-c5706237-32e2-4d38-afa4-7f4a36efd9b7-48</value>
  </correctResponse>
  </responseDeclaration>
- <itemBody>
  <p>Identify the missing word in this famous quote from Shakespeare's Richard III.</p>
  - <p>
    Now is the winter of our discontent Made glorious summer by this sun of
    - <inlineChoiceInteraction responseIdentifier="RD-d81acfa0-3f41-4da9-91dc-de9c7b2d6920-46" shuffle="true">
      <inlineChoice identifier="IC-139069b0-7960-4f4d-a39e-c9d5a28cb045-45">Gloucester</inlineChoice>
      <inlineChoice identifier="IC-fe94790b-5e62-470a-a260-cfc05c98b05e-47">Lancaster</inlineChoice>
      <inlineChoice identifier="IC-c5706237-32e2-4d38-afa4-7f4a36efd9b7-48">York</inlineChoice>
    </inlineChoiceInteraction>
    ; And all the clouds that lour'd upon our house In the deep bosom of the ocean buried.
  </p>
  </itemBody>
</assessmentItem>
```

Fig. 23.1 A QTI item represented in XML

23.3 Specifying an Assessment with a QTI Authoring Tool

While the QTI item represented in XML is easy for a computer to interpret and render, it is difficult or even impossible for practitioners to understand and specify the QTI item. The QTI authoring tool is expected to provide them with a user-friendly authoring tool that creates this XML.

However, it is not easy to support practitioners with a user-friendly authoring tool, because QTI is a technical specification which is designed for technical experts and computer. On the one hand, QTI contains many technical concepts (e.g., data types, interactions, and response variables) that are unfamiliar to the ordinary teachers. On the other hand, the terms and concepts used by ordinary teachers are not explicitly used in QTI. For example, some types of questions (e.g., multiple-choice, multiple-response, yes/no, and Likert) are not explicitly modelled in QTI v2, but all these types of questions can be represented using the same interaction type “simple-choice”. These different types of questions are distinguished by specifying the attributes and inner elements of the “simple-choice” interaction differently.

The QTI authoring tool should not expose ordinary teachers to this kind of technical complexities.

In addition, the requirement to integrate QTI items into a UOL makes it more difficult to design and implement the target QTI authoring tool. According to the solution suggested in QTI v2 to integrate a QTI item into a UOL, the UOL should contain a resource with a new resource type “imsqti”. The resource has to refer to a XML file representing a QTI item. The QTI item should define an outcome variable which is coupled with a property modelled in the UOL in a way that the identifier of a property is a combination of the identifier of the QTI item and the identifier of the outcome variable. It is difficult for ordinary teachers to define and manage the relationships between QTI outcome variables and LD properties. Moreover, if it is needed to define a questionnaire with more than one question, it will be not convenient to model more than one resource with corresponding QTI items and to couple associated outcome variables and properties.

In order to support ordinary teachers to author questions conforming to the QTI v2 and integrate them into UOLs, we developed an open source, platform-independent QTI authoring tool. We tried to provide ordinary practitioners with an easy way of creating and editing a core set of the more commonly used question types. The user interface is designed by using the concepts with which ordinary practitioners are familiar (e.g., multiple-choice question and correct answer), rather than directly using the technical concepts (e.g., simple-choice interaction and response variable). While the user edits questions in a user interface appropriate to the practitioners, the authoring tool will change the format of the defined questions and generate XML code according to QTI v2. Thus, acceptability and productivity

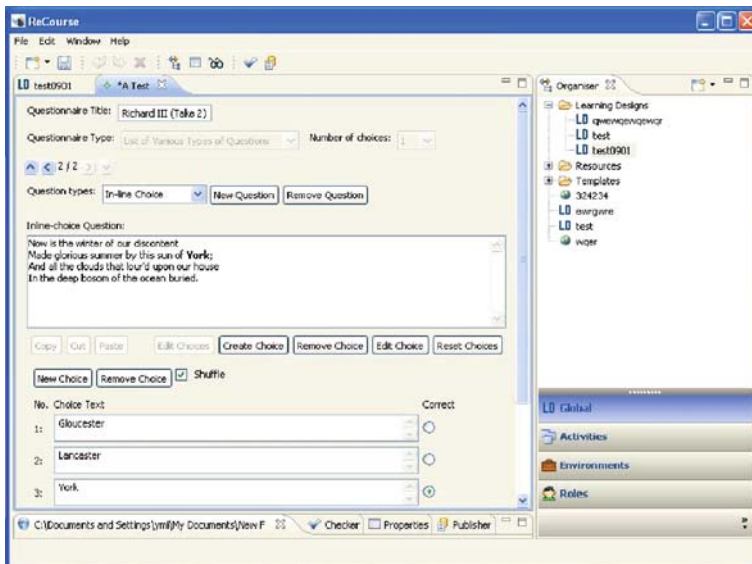


Fig. 23.2 The user interface of the QTI authoring tool integrated with ReCourse

can be greatly increased. Figure 23.2 shows the user interface of the QTI authoring tool. This sub-section describes the procedure to create and edit the assessment.

23.3.1 Creating a Questionnaire

As mentioned above, the suggested solution to integrate QTI and LD is restricted to embed the QTI item in a UOL. It is not supported to integrate a QTI test (which contains a set of hierarchically structured items) as a single resource of a UOL. As a consequence, the teachers cannot easily to arrange a test with multiple questions in a teaching/learning process as a formative assessment or a summative assessment.

In order to solve the problem, we intentionally misuse the concept of QTI item when developing the QTI authoring tool. In QTI, an item can be used to model a complicated question which combines multiple interactions with various types. According to our solution, a questionnaire with a list of questions can be modelled as a single QTI item that contains a list of interactions. Each interaction is used to model a simple question. For the user of the QTI authoring tool, it is not necessary to know the concepts of item and interaction. The authoring tool enables to create a questionnaire which consists of a sequence of questions with various question types. Currently, the QTI authoring tool can support several teacher-familiar and commonly-used question types: multiple-choice, multiple-response, yes/no, Likert, fill-in-the-blank, essay, and inline-choice. It is planned to include more question types such as matching, ordering, and slider.

The authoring process in our QTI authoring tool is not driven by the QTI information model. Using the tool, the user has to create/edit/save a questionnaire as an assessment resource. A questionnaire usually consists of the same question type, for example, a list of multiple-choice questions. In order to support the creation of a list of questions of the same type, the user can select the questionnaire type, for example, a multiple-choice or Likert. If a questionnaire type is selected, all questions in this questionnaire will be automatically defined as the selected type. Such a design increases the efficiency to create the questionnaire with a list of questions of the same type. If the questions within a questionnaire have different types, a questionnaire type “mixed” can be chosen. Then, the user has to choose a question type when creating a question within the questionnaire. In addition, if all questions in a questionnaire not only have the same type (e.g., multiple-choice, multiple-response, yes/no, Likert, and inline-choice), but also have the same number of choices (e.g., four choices or five choices), the user can assign the attribute “number of choices” a number ranged from one to ten. If so, whenever a new question is created, a given number of choices will be created automatically. In the navigation tool bar there are two numbers in between four arrows. The second one indicates the total number of questions of the questionnaire and the first one indicates which question is currently editable in the queue. The left-arrow and the right-arrow are used to display the preceding question or the succeeding question in the queue, respectively. The up-arrow and the down-arrow are used to show the first question or the last question. As mentioned, if all questions have the same type, it is not needed to select

question type when creating a new question. If the field “Questionnaire type” is set as “mixed”, the field “Question types” will be editable, so that the user can select a question type or an information block. When the user presses the “New Question” button, a new question will be inserted follow the currently edited question. Clicking the “Remove Question” button will result in the deletion of the currently edited question and display the following question if there is one, or the previous one if it exists.

23.3.2 Editing a Question

The authoring of questions is the main part of the authoring tool. A question comprises several aspects that can be edited. This amount of information can overwhelm the practitioners, so the authoring tool deals with this problem by offering a more comfortable and friendlier user interface based upon the type of question. For example, Fig. 23.2 shows the user interface for editing an inline-choice question. In the first phase, the user can edit a text. During this time, copy/cut/paste can be used for editing the text conveniently. When finishing the edition of the text, the user can press the “Edit Choices” button to change the phase. In the second phase, the user can select a word or a phrase, for instance, “York” in this example. Under this condition, an inline-choice can be created by pressing the “Create Choice” button. Then the edit area will automatically create the first choice and specify the marked selection as the correct answer. The tool also enables the user to create/remove new choices and change the correct answer. When changing the correct answer, the marked selection in the text will be changed accordingly. The “Shuffle” attribute is used to specify whether all choices should be rendered for the students in this order or not. The user can remove a defined inline-choice by clicking the inline-choice indicated in bold and then pressing the “Remove Choice” button. If the user wants to modify a defined inline-choice, he or she can click the inline-choice indicated in bold and then press the “Edit Choice” button. All defined choices will be displayed and can be edited further. Note that the defined question as shown in Fig. 23.2 will be automatically transformed into a XML file as shown in Fig. 23.1. The edition of other types of questions is supported in a similar way.

23.4 Implementation and Integration with ReCourse

As mentioned before, QTI v2 is an extensive and sophisticated specification that allows the creation of diverse and complex assessment. It is difficult and unnecessary to develop a full-functioned QTI authoring tool for ordinary practitioners, because they have difficulties to understand and handle most technical issues. In addition, the main aim of development of our QTI authoring tool is to support integration of QTI items with UOLs. Therefore, we need to analyse and prioritise the

components of the specification in order to provide maximum benefit for the IMS LD community with minimum cost.

23.4.1 Implementation

Our QTI authoring tool is implemented as a JAVA Eclipse plug-in. The tool can be integrated into ReCourse and a high-level assessment process modelling tool, which is described in Chap. 22. Of course, it can be extended as an independent application as well. The source code is stored in a SourceForge CVS repository. The details are:

```
Host: tencompetence.cvs.sourceforge.net
Repository path: /cvsroot/tencompetence
Connection type: pserver
User: anonymous
Module: wp6/org.tencompetence.qtieditor
```

There are a total of 21 interaction types specified in QTI v2. We plan to support the authoring of some interaction types that can be rendered by a QTI engine APIS (APIS), because only APIS is currently integrated with our LD engine CopperCore (Vogten 2004, 2006). We have implemented authoring functions for editing simple-choice-interaction, extended-text-interaction, text-entry-interaction, and inline-choice-interaction. We plan to extend the authoring functions for editing match-interaction, order-interaction, and slider-interaction in the future. As described above, rather than using concepts of item and interaction, we use the concepts of questionnaire and question in the design of the user interface. Ordinary practitioners can create and edit commonly used question types without the need to specify the technical details. The tool will transform each question type into an appropriate interaction type, and recognise the question types based on the feature of the interaction definition. For example, a questionnaire with a list of multiple-choice questions (with four choices) will be recognised through analysing the definition of all questions in the questionnaire, because the concept of multiple-choice question is not explicitly specified in QTI v2.

In this version of the tool, the response variables are created automatically as the user creates a question. However, the outcome variables and response processing are defined based on the predefined algorithm templates offered by QTI v2. The user can not define customised methods to process responses and calculate the final score. We plan to develop various grading algorithms which can be selected and the user just needs to provide a set of values to instantiate the template. It is important to note that scoring rules do not specify the internal details of particular scoring algorithms. Instead, the type of algorithm is specified and the necessary parameters are supplied.

Because our focus is on demonstrating the feasibility and usability of the QTI authoring tool on supporting the integration of assessment with learning processes,

our authoring tool currently only provides support for modelling questions in plain text. The use of audio, video and other objects in the question are not supported.

23.4.2 Integration

The QTI authoring tool has been integrated with ReCourse. ReCourse offers user interface to create/open a new-type (specified as “imsqti”) of resource. This resource can be referred by a LD item defined as an information unit in a UOL. After an “imsqti” resource has been created/opened successfully in ReCourse, the QTI authoring tool will be invoked and presented to the user. The user can create/edit a list of questions in the procedure described above. After finishing the authoring, the user can save it and then the QTI authoring tool will generate XML code according to the definition and ReCourse will manage it as a resource. When the user finishes the definition of the whole UOL, in which the questions are used, ReCourse will wrap the QTI XML file within the UOL package. Finally, a QTI resource embedded UOL will be published, populated and played like a normal UOL.

Using current version of the tool, the user does not need to manage the couplings between LD properties and QTI variables, and the authoring tool will handle it in a simple manner. Because we tried to hide the technical complexities from ordinary teachers, the flexibility provided by QTI is restricted. We plan to provide functions for experienced users to manage the couplings. It remains to be seen whether such functions can be used by certain users to develop complicated, integrated teaching, learning, and assessment processes.

References

- Almond, R. G., Steinberg, L., Mislevy, R. J.: A sample assessment using the four process framework. CSE Report 543. (Center for study of evaluation, University of California, Los Angeles 2001)
- Biggs, J. B.: Teaching for quality learning at University. *Society for Research in Higher Education* (Open University Press, Buckingham 1999)
- IMS Global Learning Consortium Inc: Learning Design Specification Version 1. <http://www.imsglobal.org/learningdesign/> (2003). Cited 21 Jun 2008
- IMS Global Learning Consortium Inc: IMS Question & Test Interoperability Version 1. <http://www.imsglobal.org/question/> (2001). Cited 1 May 2008
- IMS Global Learning Consortium Inc: IMS Question & Test Interoperability Version 2.0. <http://www.imsglobal.org/question/> (2006). Cited 1 May 2008
- Koper, R.: Modelling Units of Study from a Pedagogical Perspective: the Pedagogical Meta-model behind EML. Educational Technology Expertise Centre Open University of the Netherlands. [http://hdl.handle.net/1820/36\(2001\)](http://hdl.handle.net/1820/36(2001)). Cited 1 May 2008
- Vogten, H., Martens, H.: CopperCore. In: SourceForge: [http://sf.net/projects/coppercore\(2004\)](http://sf.net/projects/coppercore(2004)). Cited 21 May 2008
- Vogten, H., Koper, R., Martens, H., Bruggen, J. v.: Using the personal competence manager as a complementary approach to IMS learning design authoring. In: *Interactive Learning Environments* 16(1), 83–100 (2008)

Conclusion of the Book

Chapter 24

Conclusion of the Book

Rob Koper, Peter Sloep, Hans Hummel, Hubert Vogten, Jan van Bruggen, Marcus Specht, and Wolfgang Greller

24.1 Requirements for Learning Networks

The aim of this book, as described in the introduction, is to help people to further develop their professional competences during their careers by using the innovative powers of new media, mobile devices, and modern Internet services. This is related to the question of how to become and stay employable as a professional, given the permanently changing knowledge and technologies, and the change of jobs and markets during their careers. We have explained that the traditional educational method is not always the most suitable for all types of learning and especially not suitable for every professional. The reasons and related requirements for Learning Networks are given in Table 24.1.

So, in this chapter we will first sum-up what we have learned and then analyse whether we already meet all the requirements or that we should focus more on some of them in the near future.

24.2 Overview of Knowledge Gained So Far

We will now look at the knowledge gained for each of the services as they are described in the Sections.

R. Koper (✉)
Centre for Learning Sciences and Technologies, Open University of the Netherlands,
Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands
e-mail: rob.koper@ou.nl

Table 24.1 Reasons and related requirements for learning networks

No	Reason	Requirement for learning networks
1	Professionals tend to be very busy people, the available time to learn outside of the context of their work and daily life is very limited.	Create a seamless integration of this new type of professional learning into work and daily life.
2	The transferability of the knowledge gained in courses that are delivered outside of the daily life and work context of the professional is questionable.	Make the learning experience more contextualised, relate it to the problems in daily work and life and adapt it to the local situation of the professional.
3	Professionals differ highly in learning needs, preferences and prior knowledge.	A method of professional development will only be efficient when it is as adaptive, and personalised as possible, taking the specific characteristics of the (adult) professional into account. In order to adapt, good assessment and placement procedures should be available on which personalised navigational advice can be based.
4	The traditional role of the 'teacher as a knowledge source', doesn't work in these highly specialised areas anymore.	Professionals should be given direct access to the source of new knowledge and innovative solutions. The best way to do this is to involve them directly in the innovation and knowledge production itself as part of their job.
5	Learning in the specific context of a company does not always increase the employability because professionals tend to have different jobs, in different sectors during their professional careers, and company training is mostly concentrated on sector or company specific competences.	Learning networks should start from the individual career perspectives and competences and not solely from the perspective of a specific job, company or sector. The assessment of competences cross domains is a major challenge.
6	It is not possible to provide and pay enough teaching and training staff (and classroom time for the professional) to keep-up with the actual growth in demand for professional development.	Involving the collective knowledge and support powers of the peers in a Learning Network is an absolute requirement to get sufficient volume. Sharing of knowledge and peer-support is a key factor.
7	Self-directed learning as not suitable for all learning needs, and many times not the most efficient way to learn.	The availability of help, guidance, planning and navigation can be an essential requirement to attain more ambitious learning goals and to make the whole process of learning more efficient. Mix informal, self-directed professional development with the more formal type of teacher-directed training.

24.2.1 Conclusions of Section I: Social Interaction in Learning Networks

1. Section I discusses ways to stimulate social interaction within the context of a specific Learning Network. Such interaction is believed to be essential for knowledge sharing between Network peers to occur. Throughout the Section, *knowledge sharing and exchange* is the central tenet; it is used as a proxy for such activities as browsing and acquainting oneself with a new subject, further developing oneself professionally (non-formally, but also formally) and even fully re-educating oneself (formally, but also non-formally). The perspective taken is that of the individual, not that of a company or institution ‘installing’ a network for knowledge management within the organisation only. Guidelines and tools are developed that help motivate these individuals to become socially active in the first place, that facilitate individuals to organise their social behaviour in topic-bound community-like structures and, finally, that help maintain genuine communities once they have arisen,
2. The individual learner’s decision to spend time on social interactions with others results from a choice he or she makes between a variety of competing activities, that each may further his or her interests. Such learners have to divide their scarce time and their question will be what activity best contributes to their specific interests and constraints. Short-term benefits and long-term benefits should be discerned. The former relate to a learner’s immediate motivation to participate in some specific interaction with one or more peers, the latter to motives that are based on a rational analysis of the ultimate effects of these kinds of interactions. *Short-term benefits* are difficult to individuate as they very much reside in the eyes of the beholder only. Some learners may be convinced that benefits accrue to particular social interactions in their specific situations, others may not. However, *game theory* predicts that, given a few forgiving demands, collaboration will emerge as long as ultimately collaboration beats non-collaboration in terms of benefits. Minimising the transaction costs of interacting socially and the costs up-front of the interaction itself, will further help sway learners to participate in interactions with their peers. *Long-term benefits* thus are the ultimate motivating factor. They should relate to increases in the quality of the knowledge exchange. This increase comes about in terms of increased learner *self-satisfaction* and *reputation*, as well as in terms of the *deepening of the understanding* that comes in the wake of telling others about issues one only just has grasped oneself. The *guidelines* that developers of Learning Networks should heed in particular, are abolishing learner anonymity, enlarging the impact of future interactions relative to present ones, and avoiding discussions on the closing down of a Learning Network. Sharing success stories of social interactions, will help further.
3. It is imperative to organise peer interactions in a Learning Network. Two kinds of connections between users maximise the value of the Network. Relatively *weak ties* between participants guarantee Network-wide knowledge exchange,

relatively *strong ties* in close-knit communities guarantee maximum knowledge elaboration. The policy advocated here is to use weak ties selectively to forge strongly-tied communities. This is done by introducing a service called *peer support in ad-hoc transient communities*. Instigated by a request for help of some Network participant, the service populates these communities with knowledgeable peers of the requester, with the aim of having them collaboratively provide help. Ad-hoc transient communities are short lived and directed very much towards a limited goal, so they are not really communities in the ordinary sense of the term. Several guidelines should be followed to guarantee their emergence. For this, two concepts need to be discerned, *knowledge dating*, meaning the teaming up of participants, and *knowledge sharing*, referring to the actual mutual exchange of knowledge between partnering learners. Knowledge dating comes about best in Networks that are heterogeneous in terms of experience and knowledgeability of the participants (within bounds), and in which learners alternate between roles of learner and, say, tutor, mentor. Knowledge sharing only thrives when learners can be held accountable for their deeds, and when the participants in an ad-hoc transient community have clear-cut goals. These guidelines will have to be implemented in tools, which to some extent will be Learning Network specific. Latent semantic analysis is a tool used to support knowledge matching, for instance, interaction policies can be implemented in the business logic of network management systems.

4. Ad-hoc transient communities will not automatically result in the emergence of genuine communities from the still relatively weak ties forged between their respective participants. These ties need to be strengthened in successive encounters, within the confines of new ad-hoc transient communities set up for other purposes, or through other means. Tools such as community profiles play an important part in this. Several guidelines pertain to the maturation of communities and their maintenance. Learners should be empowered to set up their own *community profile* with levels of privacy of their own choosing, they should be helped to classify for themselves what they have learnt, and knowledge co-construction between them should be facilitated. Importantly, *trust formation* between them should be actively pursued through the promotion of off-task activities and reputation flagging. These guidelines have to be translated in tools that, again will to some extent be specific to the situation of a specific Learning Network and its inhabitants. It seems plausible that a tool for personal profiling and can piggy-back on a tool for managing personal competence development plans, which in turn can also be used to store data on actual progress in competence development over time. This then would provide a sufficiently detailed profile of some learner for trust formation to latch onto.

24.2.2 *Conclusions of Section II: Navigation Services for Learning Networks*

1. Section II contains guidelines on suitable techniques, criteria and steps to follow when applying a recommendation strategy for navigation advice. A navigation advice is an advice that suggest the best next step to take for a learner, given the learning needs and the history of the learner. Our R&D results have demonstrated that the so-called *hybrid recommendation strategies* are the most successful ones in terms of effectiveness, efficiency and appreciation, as can be perceived from both an educational and technological perspective. Such hybrid approaches apply various combinations of information-based and social-based recommendation techniques into recommendation strategies, combining the best of best worlds and limiting their constraints at the same time. Optimal combinations follows certain heuristic (or even algorithmic) recommendation rules that depend on the information available (e.g., the pedagogic context, the VLE being used), the type of learning (e.g., specific user group, formal or informal learning, self-directed professional development or more formal type of teacher-directed training) and the aims of the navigation service. This articulation of such an optimal *recommendation strategy* (and the recommendation algorithm it works on) as an outcome of innovative learning technology was hardly been applied to education so far, but appears powerful and promising to address some of the practical problems that practitioners now face when using either type of technique alone. Recommendation strategies therefore should be elaborated further as part of the solution for requirement 7 (see Table 24.1).
2. In the context of *personalisation* of Navigation Services, it will often be needed to retrieve information about individual learners and learning activities. Professionals differ highly in learning needs, preferences and prior knowledge. *Personalised Recommender Systems*, combining top-down and bottom-up information as an outcome of innovative learning technology was hardly used in educational contexts so far, but appeared promising for a large range of learning contexts, and could provide a good part of the solution for requirement 3 (see Table 24.1). Concentrating on user models, ontologies and metadata runs the risk of turning into a labour-intensive and hard to maintain effort, and sometimes will be hard to do in advance because of ‘open corpus’ problems of many Learning Networks. Nevertheless, some learning contexts will need this type of top-down information. We found that when personalised advice has added value, collaborative filtering of user behaviour could yield much (if not all) required information without humans actually having to intervene. Chapter 8 describes some techniques for *effortless and unobtrusive retrieval of information*. For instance, the pragmatics-based approach of tagging learning activities with learner information seems promising and works as follows: Each time a learner interacts with

a learning activity, the (current) learner model of that learner is attached to the learning activity.

3. Besides innovative learning technology output that is up-and-ready for use in Learning Networks, the designers and developers of Learning Networks will often need to customise or adapt existing technology to meet the specific aims or constraints of their networks. As an output of our experience we have presented some valuable guidelines or methods when setting up Navigation Services. Carrying out *Simulation studies* has proven to be an efficient means to determine the design of most suitable Navigation Services, and have lead to innovative technology output. Such studies can be executed without having to bother about practical real-life constraints of experimentation. Chapter 8 contains information (and concrete experimentation) about what we found to be the *steps* to be taken in such simulation studies, and the *parameters* to be considered when setting up a simulation study. Simulation studies have for instance demonstrated and lead to the innovative technology output of rating-based recommendations as a good light-weight alternative for intensive data maintenance approaches needed by ontology-based recommendations. We found that rating-based recommendations, when compared to peer-based recommendations, result in more, more satisfied, and faster goal achievement. It seems worthwhile to continue the investigation of other parameters in navigation through simulation studies, that will probably yield more innovative technology output. Iterations between simulations, system development, and field experiments provide valuable insights into the technical infrastructure needed to provide sound personalised recommendations to professionals in Learning Networks, addressing both requirements 3 and 7.
4. Eventually innovative technology output like navigation services have to proof themselves in the context of real Learning Networks. They are in need of a check-lists for setting them up, usability criteria for its use, and an evaluation framework to be properly assessed (Does the educational technology actually work, i.e. make the whole process of learning more efficient?). Although these results cannot be considered as technology output as such (in the meaning of artefacts), their availability is invaluable to arrive at innovative technology output. Our experience has shown that carrying out *evidence-based research* in Learning Networks is a quite challenging but feasible endeavour to arrive at validated technological innovations. Chapter 8 contains guidelines about which aspects need to be considered when carrying out such real-life studies. We have provided guidelines on the *design criteria*, on the *evaluation criteria*, on available *recommendation techniques* to be used, and on scientific methods and statistical techniques that can be used in such studies. Clearly these studies have suffered from limitations and need to be continued in various other setting and domains for further validation of the (promising) empirical findings so far.
5. Like with the hybrid approach towards Navigation Services (point 1), there also is a trade-off in using *Learning Technology specifications*. Current specifications that are most relevant for Navigation Services (e-portfolio, competence, learning path) already serve exchangeability well, but seem to lack sufficient meaning to

cater for personalised advice in response to requirement 3. Additional attributes have been proposed for such specifications to make them more meaningful and usable for personalised navigation purposes, without the risk of making them too hard to use and maintain. Especially the *learning path description* we propose is a good example of looking for such a light-weight (relatively small set of meta-data) but yet powerful compromise in learning technology. Describing learning paths in a formal interoperable way has several advantages. It could be argued that a suitable learning path can also be provided through recommender systems, but we feel it is essential that in selecting a particular learning path, learners have some notion of available options and variability.

24.2.3 Conclusions of Section III: Assessment and Placement Services in Learning Networks

Section III presented work on *two services* that are required in Learning Networks: assessment services that are needed to measure competence development and placement services that are needed to put new members of a learning network on an appropriate path towards competence development.

1. Assessment services for Learning Networks need to support types of assessment suited for competences and competence development and they need to be interoperable. The dominant standard for interoperability of assessments, IMS QTI, is not sufficient for competence assessments, because it provides insufficient support for modelling control-flow and information-flow. Building on the conceptual Ou-Cito model an Assessment Process Modelling Language is proposed that is activity- and process centred. *We demonstrated that APML can model complex forms of assessment* and that APML can be executed in an LD and QTI-compatible run-time environment. Further work is on its way to test whether other complex forms of assessments can be specified using APML.
2. A placement service is designed and developed that *assesses a student portfolio to see whether the outcomes of planned learning activities have already been met*. This is based on the perspective of Accreditation of Prior Learning (APL). The placement service described in Chap. 13 works on an absolute minimum set of assumptions regarding data that describe the competences of the professional as well as the outcomes of planned learning activities. We use latent semantic analysis (LSA) to compare the contents of the portfolio to learning material assuming that similarities here are indicators of similar learning outcomes. Tools and techniques for content-based placement using LSA were presented and a scenario was presented to validate such a placement service. Our current experience leads to the following conclusions:
 - (α) There is evidence that the assumption behind content-based placement is valid: similarity between portfolio content and learning material are interpreted by

experts as an indicator of similarity of learning outcomes. Ratings coincide with exemptions granted.

- (β) An important lesson learned is that to ensure that an LSA-based placement service can mimic expert decisions we need to establish first its sensitivity in the domain and then its reliability and validity as the initial results from a field trial demonstrated.

24.2.4 Conclusions of Section IV: Contextualized Learning Network Services

Personalisation and Contextualisation are complementary concepts, as Learning Networks are embedded in everyday activities, the personal access to information connected to the current context of use is an essential component of successful implementation of authentic learning networks. Successful approaches to adaptation in Learning Networks should consider personal user information in context. Contexts can have many facets and basically several best practices show how to combine personal information about the user and contextual information.

- *Designing authentic learning environments involves the creation of complex, situated and authentic task contexts as also the integration of user generated media and metadata.* The integration of remote locations seems a successful pattern for authenticity in learning networks. Approaches to enrich the classroom can also be applied to enriching remote and mobile learning network access. Key issues here are authentic media, instant media creation possibilities, enabling of network participants to situated access to information, and the connection of learning network members across contexts.
- Mobile social software provides a wide variety of possibilities for designing new services for mobile access to learning networks (enable active construction of knowledge, remote access to authentic problems, foster the reception of multiple perspectives on subject matter, allow for reflection about own knowledge, enable learning by social interaction within distributed learning networks). As in professional development the parallel development of different competences is the given reality, technology has to support the *instant and parallel access to necessary resources and learning network services*. Services can be classified and designed according to a given reference model presented in Chap. 15. Successful best practices and service types can focus on different purposes, i.e. supporting communication in learning networks, sharing content and knowledge, facilitate discussion and brainstorming, foster social awareness, guide communication and exploration in context, foster engagement and immersion.
- Peripheral Information is essential for supporting informal learning processes in learning networks, it can stimulate reflection, social awareness, engagement and interest in informal learning. *Peripheral information* is information that is independent from learning objectives, but allows learners to relate their actions to

the group activities. The important aspect of peripheral information is that it *creates a relational tension in the learners' perception, which allows the learners to relate, maintain, or to assess their actions in relation to those of their peers.* Mobile access to peripheral information even more supports reflection in action based on Schön's paradigm. Allowing the user to reflect about and be aware of the ongoing activities in the learning network on a meta-level are essential for gaining an understanding of the own learning process and placement in the learning network.

24.2.5 Conclusions of Section V: Learning Networks Integrated

One of the outcomes presented in Section V is the Integrated Learning Network infrastructure, which supports the core domain entities of a Learning Network through the provision of a collection of integrated web services. These services are the backbone for applications that support all kinds of professional development, whether this is achieving a start qualification, becoming more mature in a profession or simply keeping up-to-date in a job. Integration refers to the idea of the conceptual integration of these core domain entities in the form of a Domain Model describing the Learning Network entities and their relationships (Chap. 18). Integration also refers to a technical integration where the domain entity services have common authentication and authorization mechanisms and API structures. However, the proposed Learning Network offers openness as well, because no assumptions are being made about a deployment model, nor is the set of services fixed or closed. We described how servers can be confederated and explained how services can be extended and added to the Learning Network infrastructure (Chap. 19). Furthermore, the Learning Network services do not make assumptions about the client applications using them, and new applications can be developed to suit the particular needs of various groups of professionals. Sharing the same services, possibly in different circumstances and settings, effectively also implies sharing the same data, which in the end results in a richer and more effective user experience. These shared data are in effect the e-portfolio of the professional and the domain entity services draw on and contribute to this e-portfolio.

For providers of Learning Network Services such as, for example, human resource departments, the full basic life cycle of professional competence development has been covered in this way. The integrated architecture ties all the aspects of professional competence development together so information flows seamlessly from one end of the system to the other. An employee in a smaller branch of the company will thus be able to participate in the same training and learning activities as staff in larger central company locations. The integrated infrastructure therefore covers the full spectrum from personally deciding what kind of training is important to your own progress on the one end, to the other end where the policy or directive

of a professional body or company requires certain training to happen and would want to monitor their success rate.

Still, it has to be said, that refinements and extensions are feasible and desirable, which will be part of our future efforts in this field. Especially in the area of accessibility we see room for future work and improvements.

24.2.6 Conclusions of Section VI: Implementation Examples

Outcomes of the work are a set of concrete software applications using the integrated infrastructure of the Learning Network services. The infrastructure that these applications run on provides shared services that constitute the technical basis for enabling ubiquitous time- and space-independent professional learning using the Internet.

Firstly, the Personal Competence Manager (Chap. 20) is a tool to set up communities of professional learners. Competence profiles and competence development plans can be authored or adapted to personal requirements and shared with others. This provides a solution to the organisational management of mobile and adaptive professional development. Additional and enhancing services such as rating and tagging allow users to mutually evaluate available learning opportunities. This new approach to quality control and self-determination gives more powers to the learners to identify the level of relevance to their personal learning context and professional career development.

The PDP tool (Chap. 21) on the other hand was conceived as a very personal application, that allows anyone to steer their own learning process through planning and recording it. In a professional development setting, this goes far beyond the traditionally disconnected seminar opportunities, because it allows people to continuously document and monitor their progress, to reflect on it and to present it, where this is of benefit.

The learning design authoring tool ReCourse (Chap. 22) supports the creation of training by professionals for professionals. As a technical implementation of the IMS Learning Design specification it generates interoperable and transferable Units of Learning, that organisations can then use to up skill their staff in whatever area they want and in an infinite variety of pedagogical approaches.

Finally, we presented the QTI assessment tool (Chap. 23), which allows, to create assessment items and tests that can be connected to the above mentioned Units of Learning to complete the learning cycle.

24.3 Analysis and Conclusions

In this paragraph we will analyse whether we already meet all the seven requirements stated in Table 24.1 or that we should focus more on some of them in the near future. Each of the Sections describes concrete guidelines or concrete working tools based on R&D work. As with every research work, also lots of suggestions were

given for further research in specific areas. We will not repeat them here, they can be found in the conclusions of each chapter. We will look now into each Section and their contributions to the 7 requirements:

Section I focusses mainly on the requirement 6 (and partly 5): stimulating to share collective knowledge and stimulating individual competence development. The guidelines prescribed in this chapter are a good starting point for setting up active professional learning communities in which people feel stimulated to share and support each other.

Section II focusses on requirements 3 (adaptation, personalisation) and 7 (providing help, guidance). The perspective taken is to automate help and guidance as much as possible by using the history of successful learning behaviours of others to provide grounded advice to new comers. This is collaborative filtering approach to recommendations is combined with a additional information about the (similarities in) personal characteristics. Work has to be done still in finding ways of effortless and unobtrusive retrieval of information, for instance, the pragmatics-based approach of tagging learning activities with learner information seems promising and works as follows: Each time a learner interacts with a learning activity, the (current) learner model of that learner is attached to the learning activity. The general approach is very promising, it can be used in any Learning Network and it provides added value in terms of navigational support without the need to use human effort to provide this advice. In complex, ever changing Learning Networks this is also not doable for humans to keep up with all the learning facilities offered and individual characteristics of the users of the Learning Network. The Section furthermore provides useful tips at the research level: how to setup and use simulations and real life experiments to evaluate the success of certain measures within the Learning Network. For researchers and students in this area this could provide very valuable information.

Section III is a counterpart of Section II: when we want more adaptation and personalisation, as a result the assessment procedures should change. This is also the case when we want to stimulate more informal and non-formal learning: how do we assess and acknowledge the outcomes of these forms of learning? As a result this Section is also focussed on requirement 3 and 7, but from another perspective: assessment and placement. The approach taken in this chapter to use language technologies as a base is a very promising one, but is still in experimental stage. Mainly because of the complexity of these language technologies (e.g. LSA), the state of the supporting software and the technical requirements for this software. On a small scale, implementations are however possible. Several areas of further work were identified within the content-based approach to placement, including issues in corpus creation, automatic semantic segmentation of documents as well as recognition of text segments that can be used as a kind of metadata: keywords, abstracts, summaries, indices and learning objectives. Finally, further work is needed to broaden the type of data (metadata, ontological descriptions) that the placement service can handle.

A bit different is the status of the Assessment Process Modelling Language (APML). Because this language can be implemented using existing learning tech-

nologies like IMS QTI and IMS LD, in principle this could be applied in a Learning Network already. Also some first experimental software implementations have been tested in practice.

Section IV is focussing on requirements 1 (seamless integration), 2 (contextualised learning) and 4 (direct access to knowledge resources). It is still highly ambitious and needs further work, but also concrete guidelines and results are already mentioned. Specifically in this area a lot of future work is required in order to provide learning opportunities integrated in the workplace and in daily life.

Sections V and VI are somewhat different in the sense that they do not focus on a specific set of requirements, but describes Models, Infrastructures and Tools for the integration of all of them. The Service approach has been described and a variety of tools have been presented. Many more tools can be envisaged, and in e.g. the TENCompetence project many more tools using this infrastructure are in development that are not mentioned in these Sections. For providers of Learning Network Services such as, for example, human resource departments, the full basic life cycle of professional competence development has been covered by using the Integrated Services described in these Sections. The integrated architecture ties all the aspects of professional competence development together so information flows seamlessly from one end of the system to the other. An employee in a smaller branch of the company will thus be able to participate in the same training and learning activities as staff in larger central company locations. The integrated infrastructure therefore covers the full spectrum from personally deciding what kind of training is important to your own progress on the one end, to the other end where the policy or directive of a professional body or company requires certain training to happen and would want to monitor their success rate. Still, it has to be said, that refinements and extensions are feasible and desirable, which will be part of our future efforts in this field. Especially in the area of accessibility there is room for future work and improvements.

To conclude these Sections and this book, we see that a lot of concrete tools and guidelines have been developed and tested already, but especially in the field of requirements 1 (seamless integration) and 2 (contextualisation) future work should provide additional solutions.

Index

A

Ability, 323
Accountability, 50–51
 continuity, 50, 51
 history, 51
 recognisability, 50, 51
Accreditation, 199
Accreditation of Prior Learning (APL), 168
Actions, 324, 327
Actor, 324
Adaptation goal, 246
Adaptation mean, 246
Adaptation target, 245
Adaptive hypermedia, 247
Adaptive instruction, 245
Ad-hoc transient community, 48, 50, 54
 Learning Network Service, 58, 59
 peer selection criteria, 61
Aggregator, 302
Anchored instruction, 284
Artefact, 184
Assessment, 172, 175, 334
Assessment construction, 179
Assessment design, 179
Assessment process modelling language, 182
Assessment run, 180
Assessment services, 168
Authentic learning, 273, 285
Authentic tasks, 285
Awareness, 291

B

Basic competences, 319
Blended learning, 44

C

Classification, 245
Cognitive modeling, 220
Cognitive realism, 275

Common vocabulary, 197
Community, 334
 intensional, 20
 why participate, 19, 21, 22, 23, 24
Community-generated content, 262
Community interaction, 262
Competence, 318, 319, 333
Competence assessment, 327
Competence description, 82
Competence development, 316, 333
Competence development plan, 333
Competence map, 327
Competence outcome, 196
Competence profile, 327, 333
The conceptual model, 316
Conceptual simulation, 118
Content-based placement, 204
Context, 240
Context-aware systems, 247
Context-enriched learning content, 263
Context-specific learning support, 257
Contextualised language learning, 261
Contextualised search filters, 262
Contextualization, 246
Control layer, 301

D

Decision making, 181
Discrimination, 206
Distributed cooperative task work, 277
Domain entity services, 331, 348
Domain model, 334
Domain ontology, 108
Dublin descriptors, 197

E

e-assessments, 172, 173
Ecological niche, 317
Educational meta-model for assessment, 177

- Effectiveness evaluation, 284
- Electronic portfolio, 202
- Emergence, 14
- Engagement and immersion, 264
- E-portfolio description, 81
- EQF levels, 327
- European Qualification Framework, 197, 320, 333, 370
- Evaluation criteria, 98
- Events, 318
- Evidence creation, 183, 184
- Excursion, 284

- F**
- Facilitating discussion and brainstorming, 261
- Field trips, 276
- Fieldwork, 280
- Formal assessment process model, 187
- Functional competence, 319

- G**
- General text parser (GTP), 213
- Guiding communication, 264

- H**
- Heterogeneity, 44, 46–48, 58
 - effect on learning, 48
 - experience, 47
 - lurker, 48
 - person characteristics, 47
 - poster, 48
 - trendsetters, 47
- HR XML, 202
- HTTP, 336
 - DELETE, 335
 - GET, 335
 - MIME type, 335
 - POST, 335
 - PUT, 335
 - response, 350
 - status code, 335, 336, 338, 350

- I**
- IMS Learning Design, 173
- Indicator, 300
- Indirect social interactions, 327
- Information-based approaches, 83
- Information matching, 83
- Interaction footprint, 301
- Interoperability, 173
- Interpersonal trust formation, 34
 - community characteristics, 35
 - first impression, 36
 - word of mouth, 37

- Item, 334
- Item construction, 178

- K**
- Killalea State Park Expeditions, 282
- Knowledge dating, 43–54, 402
 - heterogeneity, 46
 - matching, 46
 - roles, 48
- Knowledge resource, 325
- Knowledge sharing, 20
 - accountability, 50
 - community, 20
 - goal, 52

- L**
- Latent semantic analysis (LSA), 214, 220
- Learner autonomy, 43, 446
- Learner context, 308
- Learning, 317
- Learning activities, 324, 334
- Learning in authentic contexts, 262
- Learning design, 328
- Learning goals, 325
- Learning Network, 325
 - online social network, 17
- Learning Network domain entity services, 332, 334
- Learning Network infrastructure, 330, 331, 332, 334, 356
- Learning networks, 250, 284, 315
- Learning Network server, 330, 338, 339
- Learning organisations, 317
- Learning path, 86
- Learning path description, 80
- Learning path specification, 81, 154
- Learning in practice, 296
- Learning teams, 317
- Level of competence, 319
- Linkages, 200

- M**
- Message, 333
- Mobile extension, 260
- Mobile language learning, 261
- Mobile social software, 258
- Model for navigation services, 86
- Multiple perspectives on real-world objects and locations, 262

- N**
- Natural Language Processing (NLP), 211
- Navigation, 152
- Navigation support, 298

O

Ontology, 203
 Open source, 329
 OUNL/CITO model, 173, 177

P

Participation, 298
 PDP, 324
 Pedagogical rules, 117
 Peripheral information, 294
 Personal development plan, 325
 Personalised Recommender Systems (PRS), 77
 Personalization, 219
 Personal trust formation, 29, 34, 36–40
 personal history, 39
 Placement, 168
 Placement service, 168
 Positioning, 203
 Prisoners' dilemma, 21
 Problem based learning, 284
 Process structure model, 183
 Professional development, 320
 Professional history, 197
 Proficiency level, 319, 333
 Profile, 120, 122, 124, 125, 130, 131, 132, 134,
 139, 140, 146
 foster interaction, 27
 Profiling, 198
 Progressive disclosure, 160

Q

QTI, 172
 QTI specification, 168

R

RDCEO, 202
 Realistic tasks, 284
 Recommendation plug-ins, 105
 Recommendation strategy, 108
 Recommendation techniques, 82
 Recommender systems, 74
 Reference model, 259–260
 Reflection, 291
 Reliability, 233
 Remote Access to Field Trips project, 275
 Resource, 334
 Response processing, 180
 REST, 334, 335, 344, 356
 representation, 335, 337
 Results, 317, 327
 Retrieval, 211
 R-framework, 214
 Run-time environment, 191

S

Scalability, 227
 Screen test, 221
 Self-categorisation, 30
 rating, 35
 tagging, 35
 Self-management, 30
 creation of contacts, 31
 creation of learning activities, 31
 creation of online learning communities, 31
 creation of profile, 30
 Self-organisation, 30
 comment, 33
 knowledge co-creation, 33
 recommend, 33
 share, 33
 Self-regulation, 30, 292
 control, 28, 30
 level of privacy, 30
 Sensitivity, 227, 232
 Sensor, 300
 Service-based Learning Design Player (SLeD),
 191
 Servlet, 330, 350, 354, 375
 Sharing content, 406
 Sharing knowledge, 19
 Simulations, 88–89
 The singular-value decomposition (SVD), 214
 Situated learning, 284, 291, 296
 Sociability, 14,
 Social affordance, 299
 Social awareness, 263
 Social-based approaches, 77
 Social interaction pattern, 50
 Social presence, 30
 Sound recommendations, 116
 Standardised extendable architecture, 266
 Stereotype filtering technique, 108
 Sydney Olympic Park project, 279

T

A technical framework for contextualized
 learning support, 266
 TENCompetence API, 329
 access rules, 334
 attained-level element, 347
 authentication, 330, 338, 364
 authorisation, 330, 339, 342
 authorisation rules, 339, 340
 business logic, 360
 canonical element, 356
 creator element, 340, 342
 default-rights element, 341

TENCompetence API, (cont.)

- discovery service, 331, 336
- link element, 336, 340
- password, 338
- personal element, 340
- registered element, 342, 346
- update-users, 341
- URL pattern, 336, 337, 338, 343, 344, 364
- user account, 338
- user id, 338

Term–document matrix, 220

Text to Matrix Generator (TMG), 214

U

Units of learning, 325

Usability testing, 282

The use case model, 314

User, 337

User model, 243

V

Validation scenario, 226

Validity, 233

Vectors, 213

Visualizations, 295

W

Way-finding, 74

Waylay, 299